AD-A152 996 PROTOCOLS AND SECURITY IN THE WWMCCS (WORLDWIDE MILITARY COMMAND AND CONT..(U) INSTITUTE FOR DEFENSE ANALYSES ALEXANDRIA VA T C BARTEE ET AL. NOV 84 1/2

UNCLASSIFIED IDA-P-1796 IDA/HQ-84-29059 MDA903-84-C-0031 F/G 17/2 NL

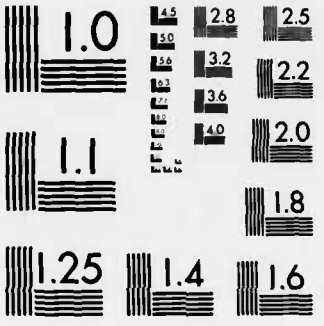

(2)

IDA PAPER P-1796

PROTOCOLS AND SECURITY IN THE WWMCCS INFORMATION SYSTEM (WIS)

T. C. Bartee
H. B. Heiden
J. M. McQuillan
S. T. Walker

November 1984

Prepared for
Office of the Under Secretary of Defense for Research and Engineering

INSTITUTE FOR DEFENSE ANALYSES
1801 N. Beauregard Street, Alexandria, VA 22311

IDA Log No. HQ 84-29059

The work reported in this document was conducted under contract MDA 903 84 C 0031 for the Department of Defense. The publication of this IDA Paper does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official position of that agency.

This paper has been reviewed by IDA to assure that it meets high standards of thoughness, objectivity, and sound analytical methodology and that the conclusions stem from the methodology. IDA does not, however, necessarily endorse the conclusions or recommendations that it may contain.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE	READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO. ADA152996 / 3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Protocols and Security in the WWMCCS Information System (WIS)	5. TYPE OF REPORT & PERIOD COVERED Final May-November 1984
	6. PERFORMING ORG. REPORT NUMBER IDA PAPER P-1796
7. AUTHOR(s) T.C. Bartee, H.B. Heiden, J.M. McQuillan, S.T. Walker	8. CONTRACT OR GRANT NUMBER(s) MDA 903 84 C 0031
9. PERFORMING ORGANIZATION NAME AND ADDRESS Institute for Defense Analyses 1801 N. Beauregard Street Alexandria, VA 22311	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Task T-2-120
11. CONTROLLING OFFICE NAME AND ADDRESS WIS/JPMO/AD Washington, D.C. 20330	12. REPORT DATE ~~February 1985~~ Nov. 84 / 13. NUMBER OF PAGES 123
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) DoD-IDA Management Office 1801 N. Beauregard Street Alexandria, VA 22311	15. SECURITY CLASS. (of this report) UNCLASSIFIED / 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE --

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

N/A

18. SUPPLEMENTARY NOTES

N/A

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

WWMCCS Information System; Military Command, Control and Communications; local area networks; network security; network architecture; network protocol; network host protocol; network front-end protocol; network processors

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Improving the survivability of the Nation's military command, control, and communications systems, under the spectrum of threats foreseen for the future, is one of the most important goals of the WWMCCS Information System (WIS) Joint Program Management Office. Critical to the survivability of the modernized WIS will be the connectivity it will provide its operational users. The closer user connectivity approaches full connectivity the greater will be the probability of survival under stress--stress being caused by destruction of

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. (Continued)
nodes, electronic warfare and/or increased traffic. The proliferation of packet-switched networks has great potential for offering the WIS a high degree of user connectivity because of the ability of packet-switching to asynchronously time-share communication links and to "Alternate-route" traffic around failed links or nodes and to bypass congested nodes.

The objective of this report is to provide guidance to the developers of the WIS architecture in achieving an acceptable level of survivable user connectivity utilizing, as much as possible, existing packet-switched networks. Particular emphasis is given to security, an important aspect of any DoD command and control system.

RE: Classified References, Distribution Unlimited
No change per Ms. Betty Pringle, IDA

IDA PAPER P-1796

PROTOCOLS AND SECURITY IN THE WWMCCS INFORMATION SYSTEM (WIS)

T. C. Bartee
H. B. Heiden
J. M. McQuillan
S. T. Walker

November 1984

INSTITUTE FOR DEFENSE ANALYSES

Contract MDA 903 84 C 0031
Task T-2-120

PREFACE

Improving the survivability of the Nation's military command, control, and communications systems, under the spectrum of threats foreseen for the future, is one of the most important goals of the WWMCCS* Information System (WIS) Joint Program Management Office. Critical to the survivability of the modernized WIS will be the connectivity it will provide its operational users. The closer user connectivity approaches full connectivity the greater will be the probability of survival under stress--and/or increased traffic. The proliferation of packet-switched networks has great potential for offering the WIS a high degree of user connectivity because of the ability of packet-switching to asynchronously time-share communication links and to "alternate-route" traffic around failed links or nodes and to bypass congested nodes.

The objective of this report is to provide guidance to the developers of the WIS architecture in achieving an acceptable level of survivable user connectivity utilizing, as much as possible, existing packet-switched networks. Particular emphasis is given to security, an important aspect of any DoD command and control system.

---

*Worldwide Military Command and Control System.

CONTENTS

ABBREVIATIONS

AC	Access Controller
ANSI	American National Standards Institute
ARPANET	Defense Advanced Research Projects Agency Computer Network
ASCII	American Standard Code for Information Interchange
BBN	Bolt, Beranek and Neumann
BIU	Basic Interface Unit
CATV	Community Antenna Television
CCITT	International Consultative Committee on Telegraphy and Telephone
CENTCOM	Central Command
CSMA/CD	Carrier Sense Multiple Access/Collision Detection
COINS	Community On-Line Intelligence System
CUS	Common User Support
DCA	Defense Communications Agency
DES	Data Encryption System
DDN	Defense Data Network
DoD	Department of Defense
DoDIIS	Department of Defense Information Systems Internat Study
EBCDIC	Extended Binary Coded Decimal Interchange Code
EUCOM	European Command
FDDI	Fiber Distributed Data Interface
FTAM	File Transfer, Access and Management
FTP	File Transfer Protocol
GCOS	General Comprehensive Operating Supervisor
HFE	Host Front End
HIP	Host to IPLI
HLDC	High-Level Data Link Control
ICMP	Internet Control Message Processor
IEEE	Institute of Electrical and Electronic Engineers
IMP	Internet Message Processor
IP	Internetwork Protocol
IPLI	Internet Private Line Interface
ISO	International Standards Organization

JDSSC	Joint Data Systems Support Center
JOPES	Joint Operational Planning System
JPMO	Joint Program Management Office
LAN	Local Area Network
LLC	Logic Link Control
MAC	Media Access Control
MAC	Message Authentication Check
MAC	Message Authentication Code
NAS	Name Authentication Server
NBS	National Bureau of Standards
NCP	Network Control Protocol
NMCC	National Military Command Center
NSA	National Security Agency
OJCS	Office of the Joint Chiefs of Staff
OS	Operating System
OSI	Open Systems Interconnection
PAD	Packet Assembler/Disassembler
PWIN	Prototype WWMCCS Intercomputer Network
RNP	Remote Network Processor
SAP	Self-Authentication Password
SAP	Service Access Protocol
SL	Security Level
SMTP	Simple Mail Transfer Protocol
TAC	Terminal Access Controller
TACACS	TAC Access Control System
TCB	Trusted Computer Base
TCP	Transmission Control Protocol
TFTP	Trivial File Transfer Protocol
TNB	Trusted Network Base
TP	Transport Protocol
ULP	Upper Level Protocol
VTP	Virtual Terminal Protocol
WIN	WWMCCS Intercomputer Network
WINCS	WWMCCS Intercomputer Network Communications System
WIS	WWMCCS Information System
WWMCCS	Worldwide Military Command and Control System

I. INTRODUCTION

A. BACKGROUND

The WWMCCS Information System (WIS) Modernization Program plans to change the transport protocol used in the WWMCCS Intercomputer Network (WIN) from the NCP protocol presently used to the DoD Standard TCP/IP protocols. The Honeywell 6000 computers currently in use will continue to be used for some time and it will therefore be necessary to use the TCP/IP protocols for these computers. In order to make this change it is necessary to evolve a management strategy for making the transition, since several important issues must be resolved before actual implementation of TCP/IP can begin.

Among the issues to be resolved are whether to place the TCP/IP protocols in the 6000 computers or to purchase front-end processors and off-load these and perhaps other network-related protocols. The possibility of using a translator or software converter between NCP and TCP/IP has also been discussed, and this possibility needs to be analyzed.

Further, plans need to be made for phasing NCP out and TCP/IP in, and the successful implementation staging of this change at various installations will depend on the use of a satisfactory basic strategy.

The WIS Modernization Program also plans the installation of local area nets (LANs) at WWMCCS sites. These LANs initially will link area work stations to the Honeywell 6000 Series computers and to the WWMCCS Intercomputer Network. As the modernization program progresses, additional resources and tasks will be added to the LANs, including some of the functions

now resident in the 6000 series computers. Only after this transfer of functions has reached its final stages can the 6000 computer series be phased out of WIS in order for the transition process to be satisfactory.

It will be necessary to provide security, including access control, authentication, and accounting in the LANs. Presently, most sites are operated at the Top Secret level and users must be cleared at that level. However, about 98 percent of the data is only Secret, or less, and so eventually it will be expedient to have a multilevel secure system. In this case the LANs must protect classified information passing through them. This report presents study results concerning the techniques for this. These techniques are now in their developmental stages. An important consideration is that while LAN security controls can initially be satisfactory for a system-high environment, it will be necessary to transition to a controlled and possibly multilevel secure system.

B. PROTOCOL STUDY OVERVIEW

Careful planning will be required to transition from the dedicated systems using the Honeywell Series 6000 computers to the adaptive architecture now planned using LANs. The Honeywell 6000s were installed starting in 1972 and so the present centralized operation at WWMCCS sites has a long period of development. In 1978 the WIN was declared operational and in 1982 the Defense Data Network (DDN) began to be used by the WIN. The WIN now uses the DDN C/30s and the circuit layouts have been rearranged for the DDN system. There are 39 Honeywell H-6000 computers and 80 Level-6 processors used as terminal concentrators and batch control devices.

In order to process information in a system with computers distributed geographically, standard protocols are used to handle the data transfers. These protocols must be arranged so

that users can pass information through and into the system computers satisfactorily. In order to economize on system design and implementation, standard protocols are generally used in the various system components. This report discusses the international standard reference model and discusses how WIS protocols relate to this standard. The discussion includes observations on current protocols and comments concerning the state of present WWMCCS systems protocols and is relative to the state-of-the-art in protocols.

There are a number of considerations concerning the protocols to be used in the LANs discussed here. The international standards organizations now support several LAN protocols, while others are in development. The status, advantages, and disadvantages of these are presented.

The above is followed by a discussion of DoD protocols, and this is followed by some of the advantages of DoD protocols to WIS.

The subject of protocol translation and the possibility of use of a protocol converter between NCP and TCP/IP is then presented and the position is taken that such a plan is not advisable. The reasons for this position are given in detail in these sections. Included are high cost, limited functionality, and lack of reliability.

The choice of whether to use network front ends or the host itself to implement the TCP/IP protocol suite is then discussed. The position is taken that use of front ends is desirable and supporting arguments are presented for that conclusion.

The final sections discuss the transition stages and give general considerations concerning how to make this transition smooth and satisfactory. The subject of a later transition from TCP to the TP protocol now in development by international organizations is also treated in some detail.

We also made a small study of the Applitek system which is being considered for use as an access system for the LANs in

WIS. The Applitek system is more complicated than most now in use, but appears to give superior operating performance with regard to cable usage. Some considerations concerning reliability, security, and system tuning are given in Section II.K.

C. LAN SECURITY ISSUES FOR WIS

In Section III of this report, the security aspects of the WIS LAN architecture are examined from the specific LAN perspective and from an overall system security vantage point. Following a brief review of the present WWMCCS system and its planned evolution in Sections III-A and III-B, an overview of the physical, communications, and computer security aspects of LANs is described in Section III-C. If trusted LANs and/or end-to-end encryption units were available today for use on the WIS LAN, they would provide an excellent framework for moving the WIS into a multilevel secure mode of operation. Unfortunately, neither Blacker for end-to-end encryption nor the various trusted LAN projects currently underway will be available in the next three to five years. Deployable LANs can probably be expected within five years and they hold the most promise for accelerating the rapid evolution of WIS in the early 1990s. In the short term, the development of Message Authentication Check (MAC) systems for ensuring the integrity of messages in transit over a physically protected LAN hold promise for providing at least limited multilevel secure operation.

Beyond the LAN itself, though, it is necessary to understand what the implications of an untrusted communications media will have on the use of trusted and untrusted hosts on WIS. Section III-D provides an analysis of the implications of operating both trusted and untrusted hosts on a simple network (consisting of just simple individual interconnecting wires) in order to gain an understanding of how a WIS LAN without trusted bus interface units or end-to-end encryption

might function. If the untrusted LAN were protected to the overall system-high level, and if MACs were used for all communications across the LAN, it may be possible to operate in a network system-high environment with all untrusted hosts at the LAN site at the system-high level and trusted hosts allowed to range downward from that system-high level, depending upon their level of trust. While this represents something less than the long-term goal for WIS, it does present a very useful environment which can be achieved in the near future with very little development effort.

If several WIS LANs must operate at different system-high levels and communications must be established between processes working at comparable levels of security, it will be necessary to call upon some form of network access control mechanisms to mediate access among these LANs. The only such mechanisms currently under development are those envisioned by Blacker for the DDN. Such facilities are expected to be available in the 1989 to 1990 time frame and should provide a suitable mechanism for this level of network access control. On the other hand, if it is necessary to operate untrusted hosts on the same LAN at different security levels, then it will be necessary to wait until trusted LANs or end-to-end encryption as applied to LANs becomes available in the early 1990s.

Section III-E presents an analysis of the options available for user authentication in a WIS LAN environment. The advantages and disadvantages of host versus work station log-in systems are compared. Several log-in mechanisms currently in operation on major networks are described for possible use on the WIS. The ideas for self-authenticating passwords as log-in mechanisms on the DDN, coupled with the work station to Name Authenticating Server notions of several commercial networks, provide a means of uniform and consistent log-in procedures with a high degree of integrity, irrespective of the type of trusted or untrusted LAN currently being employed.

It is essential, as the WIS program selects and begins to install its LAN and related security and access control mechanisms, that it be in a position to maximize the use of trusted host computer systems and rapidly evolving trusted LAN and access control mechanisms. A framework for understanding the major issues confronting this evolution is presented in Section III of this report.

II. PROTOCOLS AND SCHEDULING

Transitioning from the self-contained, dedicated systems of the WIN to the adaptive architecture of WIS computers and work stations on a local area net (LAN) without disrupting secure ongoing operational support to command and control users will require careful planning, system engineering, and execution. At the heart of this problem is the transition from the current unique WIN protocol suite to WIS protocols, for it is these protocols that make things happen in data processing exchanges.

Simply stated, protocols are a set of formal conventions or procedures that make possible information exchanges between logical processes running on computers, and between computers and terminals or peripheral devices. The burgeoning requirements for dissimilar computers to exchange information has fostered the development of an open (nonproprietary) hierarchical or layered protocol structure, where each layer performs certain functions as independent as possible of the other layers. A perfect set of protocols would permit system users to attach any computer or work station into any type of network like one can plug any lamp into any wall socket. Unfortunately, there will probably never be a perfect universal protocol suite; and complete industry-wide standardization, while being pursued, is still a distant prospect.

The rapid evolution of intercomputer communications has also left many unanswered questions in the area of security. Systems designed to be flexible, resilient, and open, as is the case with most local area nets and the WIS, are much more difficult to secure than closed systems like WIN. As with protocols, perfection is not likely. A perfectly secure system that

is also cost effective, user friendly, and can provide for the timely acquisition, storage processing, transmission, and display of critical command and control information, is not yet in the offing.

The problem, then, is not only the difficult transition from one system architecture to another without disrupting operations. It is also being able to make this transition in an ever-changing commercial and technological environment without getting painted into a corner or stuck with obsolete products or an unsecurable system.

A. BACKGROUND

Installation of the large Honeywell 6000 computers at major command headquarters began in 1972. Centralized operational software support came from DCA, but because the H-6000s were originally installed as standalone systems, owned and operated by their respective command headquarters, specialized applications and support software was locally developed to support site-unique functions. In 1975 the Prototype WWMCCS Intercomputer Network (PWIN) was born out of experiments conducted to link three H-6000 sites together using modified ARPANET hardware and software. System software was developed to support file transfers and a capability to connect remote terminals evolved.

The H-6000 was connected to the network via a Honeywell Datanet 355 processor using a modified ARPANET Network Control Protocol (NCP). Terminals were also connected to the DN-355. In 1978 the WWMCCS Intercomputer Network (WIN) was declared an operational system, as all (20) WWMCCS sites were successfully interconnected via 15 switching nodes and associated communications circuits. Improvements in computer and network hardware and software continued. Existing applications programs were refined and new ones added. The H-6000 host operating system was improved for higher throughput. DN-6661 and DN-6678

communications processors replaced the DN-355 as the H-6000 network interface. In 1982 the communications subsystem was brought into the fold of the Defense Data Network (DDN). All packet switches were replaced with newer DDN C/30s and the circuit topology was completely restructured.

1. The WWMCCS Intercomputer Network (WIN)

The present Worldwide Military Command and Control System (WWMCCS) consists of thirty-nine H-6000 computer hosts and approximately eighty Honeywell Level 6 processors. The Level 6 processors act as terminal concentrators and access control devices, and are connected directly to H-6000 hosts over dedicated circuits. Communications between the H-6000 computer hosts is provided exclusively by the WIN Communications System (WINCS). At any given moment, twenty-five of the thirty-nine H-6000 computed hosts are interconnected by the WINCS.

2. Open System Architecture

The term "network architecture" is commonly used to describe how hardware components of a communications system are tied together to provide the functionality of the system. This type of network description tells us little about how a data system functions, or the way information is exchanged and controlled. More than anything else, the protocol structure used for processing and transmitting information determines how and how well a system may perform.

Procedures used to formalize the transfer of information between computers and between terminals and computers were developed as data communications developed. Monolithic protocols were developed by mainframe manufacturers. Most were proprietary, vendor-unique protocols which would work only with one vendor's equipment. Large computer manufacturers maintained a natural single-source advantage from the lack of industry-wide standards. But, as computers became less expensive, the notion of distributed processing between dissimilar systems became more popular with user groups.

The International Standards Organization (ISO) is a voluntary nontreaty group made up of the principal standardization body of each participating nation. ANSI represents the United States at ISO functions. In 1977 the ISO created a committee for "Open Systems Interconnection" (OSI). Rather than starting immediately to develop nonproprietary, general use standards, the committee began the framework for an open system protocol architecture that in 1979 became the seven-layer OSI reference model.

The goal of the open system architecture is intercomputer exchanges between any host or terminal over any type of network. The obvious advantage of an open system to the user is interoperability without dependency upon one vendor. The disadvantage is not as obvious. To implement and manage an open network requires that users make prudent selections from a wide variety of network resources that are optimized for specific functions. What flexibility buys is paid for by the need to be very involved and technically competent. Such is the case with the WIS. System planning, engineering, and management all become more important to the development of an open WIS architecture capable of continuous evolution.

The OSI reference model, and especially the notion of seven layers, has been universally accepted as a reference and will be used here in discussing the functions protocols must perform in any system. A specific discussion on protocol standards applicable to WIS requirements will follow a general description of the OSI seven-layer model.

It should be noted that the ISO OSI reference model is not a standard. Standards are developed by other organizations such as the Department of Defense (DoD), the American National Standards Institute (ANSI is the U.S. voice in ISO), the Institute of Electrical and Electronic Engineers (IEEE), the National Bureau of Standards (NBS), and the International Consultative Committee on Telegraphy and Telephone (CCITT), to name a few.

Manufacturers, trade organizations, individuals (as those in the IEEE), and governments all play a part in standards issued by these organizations. Also, the OSI reference model does not tell how or where protocols will be implemented. Protocols may reside in hardware, special device drivers, the operating system, etc., and still be compliant with the OSI model. Lastly, the OSI model is somewhat dated, even though it is only five years old. While the model can be used primarily for local area and long-haul networks, it was designed with the latter in mind. This focus also slighted the internetting and catanet environments primarily because these areas have developed faster than anticipated, while the international acceptance of the framework (and standards which follow) has developed as slowly as expected. Even so, the model is still an appropriate framework for all data networks.

B. THE BASIC REFERENCE MODEL FOR OPEN SYSTEMS INTERCONNECTION

In the OSI model each "layer" performs a set of closely related functions. No layer is completely dependent on any other layer, and all layers do not have to be used to perform a data transaction. The seven layers and their related functions, as they apply to WIS, are discussed in detail below and summarized in Table II-1. Many existing protocols were developed before the notion of seven layers existed, and therefore the services may not correlate exactly to the OSI reference model. References to WIS-FD-100 or the WIS Functional Description are to the 27 July 1984 final draft of the Functional Description for LAN Subsystems.

1. The Physical Layer

There is no precise definition of what the physical layer includes, but it is clear that it performs the four major functions of mechanical, electrical, functional, and procedural characteristics to activate, maintain, and deactivate physical

TABLE II-1. LAYERS AND FUNCTIONS OF THE OSI REFERENCE MODEL

OSI Layer	Functions
7. Application	- Log-on, password checks & file access - Graphics, videotext - Data/information/job/file transfer - Cost allocation (where applicable) - Synchronization & error recovery - Message handling
6. Presentation	- Session initiation/termination - Test compression/transformation - Encryption - Virtual terminal/device handlers - Code/format conversions
5. Session	- Establishes/releases process-to-process connections - Control/synchronization of data exchange - Exception reporting
4. Transport	- Reliable transfer of data - Error detection/recovery - Establishes/closes host-to-host connections - Provides host-to-host flow control - Establishes data unit size - Hides communications protocol layers from host processing
3. Network	- Sets up intranetwork routing - Packetizes data for transfer - Intranetwork status messages - Intranetwork flow control - Intranetwork addressing - Internetworking (e.g., IP)
2. Data Link	- Activates/deactivates link connections - Adds beginning and end of message indicators - Allocates channels in LAN
1. Physical	- Provides electrical interface - Handles CSMA collision detection

connections. The most popular is the older (1969), 25-pin serial and parallel connection described by EIA RS-232-C. Other popular serial interfaces are the CCITT X.20 and X.21, EIA RS-449 and MIL-STD 188-114. The most popular interfaces for today's larger systems are RS-449 and X.21, the latter being a subset of X.25. WIS-FD-100 (draft) allows either the MIL-STD or RS-449 interfaces, which are plug-compatible with the older RS-232-C, but does not allow the physical layer protocol (X.21) of the ubiquitous X.25. The choice of RS-449 over X.21 is prudent. While X.21 would adequately satisfy today's requirements, it has limited potential without major modification. It cannot pass control information simultaneously during data transfer, does not have separate send and receive timing circuits, does not have provisions for loop-back testing, frequency, or signalling rate selection, etc.

2. The Data Link Layer

The purpose of the data link layer is to provide the means to establish, maintain, and release data links between network entities.

When this level was conceived it was oriented primarily to long-haul networks. In 1980 the IEEE Computer Society standardized the IEEE-802 series as a means of connecting computer equipment to LANs. The IEEE used the ISO reference model, combining their logical link and media access control procedures in ISO level 2 and dispensing with the network level. Both the IEEE LAN protocols and the network data link protocols are applicable to the WIS LAN.

There are two types of data link protocols--byte- (or character-) oriented procedures and bit-oriented procedures. The former are as old as data communications and are widely implemented, which is why they will be with us well into the future, despite their inflexibility, half-duplex nature, and orientation to batch processing. Bit-oriented procedures were

designed more to support transaction-oriented (interactive) operations, are general in scope, and flexible in possible applications. Commonly known procedures in this category are ADCCP, HDLC, BDLC, SDLC, and UDLC. ANSI approved ADCCP as a standard in 1979, but it never gained wide acceptance. Instead, ISO sanctioned the simpler High-Level Data Link Control (HDLC) which has been adopted and embodied in X.25 by FED-STD 1041 using a balanced, asynchronous, two-way, simultaneous data transfer (LAP-B). BDLC, SDLC, and UDLC are examples of popular proprietary standards, while ADCCP, HDLC, and LAP-B are non-proprietary standards embracing a larger view of the market-place. HDLC and LAP-B are specified in the draft WIS-FD-100 for the external gateway level 2 protocol. These are both good choices, but with two cautions concerning the gateway implementation. The applicable standard for LAP-B is FED-STD 1041; the applicable specification is Appendix J of BBN Report 1822. Unfortunately, the BBN implementation of LAP-B does not match their specification, and it is to a BBN device that the IU will interface. The draft FD also calls for the BBN HDH interface, which is a host-level adaptation of HDLC. BBN is having problems with this protocol, and available documentation may not accurately describe the "soon-to-be" corrected version of HDH.

The IEEE-802 protocol series for LAN applications is essentially divided into two groups. The Logical Link Control protocol (LLC) defined by IEEE-802.2 and the Media Access Control (MAC) protocols defined by IEEE-802.3 through 802.6. LLC is used in conjunction with any of the four MAC protocols. The Logical Link Control protocol is functionally very similar to HDLC, except for a frame structure that differs from any present ISO or ANSI procedure. IEEE-802.2 is planned to be used on the LAN side of the WIS IU, linking TCP/IP with the Media Access protocol. It is also used in the bridges and LAN side of the gateways. It is a proper choice.

Three of the four Media Access Control protocols will be discussed. Not considered is IEEE 802.6. This is a project under study to put together a metropolitan local area net.

IEEE-802.3 defines a baseband or broadband bus utilizing carrier sense multiple access with collision detection (CSMA/CD). It is essentially the same as the Ethernet standard advocated by DEC, Intel, and Xerox. As presently defined, the bus can be either broadband or baseband running at 10 Mbps on coaxial cable. Work is in progress for specifications on 1, 2, and 5 Mbps standards, and for 20 Mbps for baseband.

The CSMA/CD access method has become very popular because it uses common CATV components and is well-suited for the lightly loaded LANs of most office environments. With CSMA/CD there is not a central controller to regulate access. All stations have independent access, or the same priority to the medium. Stations must listen to the data channel. If it is busy, the station is prevented from transmitting until it is clear. If no carrier is detected, the station may send. A collision will occur if two stations send at the same time. The collision will be detected, the network jammed, and both stations will back off. After a collision, all stations are required to wait a random period before attempting to transmit. The CSMA/CD requires about 500 lines of code (Pascal) to implement.

IEEE-802.4 defines a bus utilizing token passing as the access method. Again, the bus may be broadband or baseband with data rates up to 20 Mbps specified for broadband and 10 Mbps for baseband carriers. The appeal for this standard has been primarily in factory automation where the advantages of a bus topology (discussed later) are combined with the capability for prioritizing traffic. Unlike CSMA/CD, token passing is deterministic in nature. That is, it is possible to calculate the maximum delay required for a station to access the medium,

as each station receives the token, and therefore the right to access, in an ordered sequence. The delay depends on the length of the network and the number of stations on it. The monitor and control functions are the most difficult in a token-bus environment, and require about 2000-3000 lines of code (in Ada) to implement. The draft WIS functional description allows for IEEE 802.4 applications in the interfaces and cable plant, but CSMA/CD is called out for future growth in the "Notes" section of the WIS-FD-100. We believe that the token-bus would provide an expensive and non-optimal solution to the WIS requirements.

IEEE-802.5 has been sponsored by IBM and was due out this Spring. It is known that IBM has had trouble implementing the token-ring approach, which this standard would describe. It is suspected that IBM wants to delay the standard until it has its token-ring perfected. That way the IBM solution would be described exactly by the standard. While the standard is believed to be fairly complete, it is not now expected to be published for another two to three years. However, the world is not waiting. ANSI Committee X3T9.5 is working on a standard that will specify a 100-Mbps token-ring architecture using optical fiber cable plant. The standard, called Fiber Distributed Data Interface (FDDI), is due out this year. Also, other than IBM, token-ring LANs are available despite the lack of a standard. One company now offers a baseband token-ring LAN with an up-to-80-Mbps capability over twin-axial or fiber-optic cable.

Token-passing actually works best on a ring topology because the token passes by each station in turn. Each station receives from and transmits to only one station. The transmission time per station is normally quick (one bit) because the full preamble can be handled by subsequent processes. Token-ring access has been implemented with the fewest lines of code (300 in Ada).

There are several token-rings in use today using TCP/IP. The largest supports over 35 VAX and PDP-11 computers and several gateways on a 10-Mbps ring. The Army's Ballistic Research Laboratory is in the process of installing a 70-Mbps token-ring. On the West Coast, Aerospace's cable plant uses a mixture of optical fiber, coax and twin-axial cable for a token-ring LAN that runs to and through five buildings. We believe it is a grave oversight to preclude token-ring technology in the WIS architecture.

3. The Network Layer (Layer 3)

This layer accomplishes the routing and switching of data between any two systems, where a system may be a network or only a host. In monolithic protocols, such as SNA or WIN NCP, this is not a discrete function. In the WIS X.25 (level 3) part of 1822 (the 96-bit ARPANET leader) and IP perform this function.

Both X.25 and BBN 1822 packetize incoming data flows, set up what appears to be a dedicated path through the transmission media, and send the packets through these virtual circuits to their network destination where they are reassembled into messages ready for delivery to the receiving host.

The IP protocol performs a distinct function and is directly linked to the subnet processes performed by X.25 and 1822, but is included in layer 3 because it provides a communications service. IP performs a datagram, or connectionless service whereby it strips the packet of its subnet layer (X.25 or 1822) and broadcasts it, with no guarantee of path or delivery, to the receiving IP. There the new subnet protocol is affixed and the packet is delivered to the receiving system. ISO is working on a standard internet protocol. It is being based on the DoD IP and will have greater address space, but probably less functionality.

X.25 may be used to interface CUS hosts and work stations, and H-6000 processors to the LAN interface unit.

WIS-FD-100 states in one place that the gateway will interface to the DDN IMPs, and in another place it states that the gateway will interface to the IPLI. We believe the WIS gateway will interface to five DDN IPLI ports using IP and 1822. Actually, a version of IP, called HIP (Host to IPLI) for host-level IP, has been specified by WIS-FD-100 for the IPLI interface. HIP is a very-slightly modified 1822, but BBN has returned to 1822 as the IPLI interface. IP will also be used in the gateways and bridges, which both access the LAN using IEEE protocols. It should be noted that some existing implementations of IP have had addressing problems when used in the Ethernet (802.3-based) environment. The nature of the problem was that the short, 4-byte, addressing space was modified to recognize 3-byte vendor numbers with the result that data flows over the LAN could be between similar (same manufacturer) hosts only.

BBN's 1822 is a mature protocol in wide use by the ARPANET community. However, it has never gained the commercial support of interface vendors. Recently, most vendors have focused on X.25, and the availability of X.25-supported interfaces and products continues to grow. As a result, the DDN X.25 has become the network layer interface of choice. BBN 1822 will continue to be supported, but probably not improved upon, until it attrites to oblivion. It is recommended that the WIS program adopt X.25 in lieu of 1822. DDN IMPs support both protocols, but the IPLI, in its developmental configuration, supports only 1822. As IPLIs are provided by the DDN, the WIS program office needs to estimate its IPLI requirements (quantities and desired interface) and so inform the DDN Program Management Office.

4. The Transport Layer (Layer 4)

The purpose of this layer is to establish and maintain reliable end-to-end, e.g., host-to-host, communication for use by the Upper Level Protocols (ULP), layers 6 through 7. Additionally, it hides or makes transparent the communications

protocols (levels 1-3) from the user's ULPs. The transport layer is required to optimize the use of available communications services to provide satisfactory performance at minimum cost. Level 4 implementations normally make only the "best effort" to deliver a packet. Reliability is guaranteed by assuming that if a packet is not delivered within a negotiated time, it never will be, and the packet is retransmitted until receipt is acknowledged.

The ISO transport layer has not yet been defined by a standard, though a very detailed draft of the National Bureau of Standards (NBS) Transport Protocol (TP) is out for review. TP is currently the top candidate for ISO acceptance and has been implemented for test by a few mainframe vendors. Overlooking the recent ballyhoo concerning this protocol, it has yet to be thoroughly wrung out and debugged--a process that took over two years and entailed a specification rewrite for TCP.

TP is not a single protocol as is TCP. As a result of compromise and at the expense of universal interoperability, TP will exist as five protocols designated as classes 0-4. The higher the class number the greater the reliability, overhead, and implementation cost. TP-4 is functionally the same as TCP and, when mature, the one that would be appropriate for DoD use. It is expected that it will be eight to ten years before TP will be an off-the-shelf product offered by most vendors. A few vendors may have commercial offerings sooner, but until the majority offer it there is little, if any, advantage to switching from TCP.

The DoD Transmission Control Protocol (TCP) performs all of the functions required by ISO layer 4 and some of the level 5 functions as well. Besides being mandatory, it is a must in an internet environment. The advantages of TCP and the eventual evolution to TP are discussed in later paragraphs.

5. The Session Layer (Layer 5)

While the transport layer sets up host-to-host connections, the session layer's job is to set up, control, and tear down computer process-to-computer process connections. These connections are called session. To do this, the session layer associates user data streams with transport layer connections and establishes session connections between peer presentation layers of two different hosts. These connections are called "sessions." The session layer performs two categories of service--a "session administration service" and "session dialogue service." The former is responsible for binding together presentation layers of two hosts, while the latter is responsible for the control of data through delimiting and synchronization. To date no standard, procedure, or specification exists for this layer. The functions of the session layer are now performed by other (existing) protocols.

6. The Presentation Layer (Layer 6)

This layer presents information to communicating application programs/files/entities in a form that preserves the meaning while resolving syntax or code differences. It also initiates and terminates session level calls, formats data, makes distributed data (as a distributed file) appear as a single data structure, and performs text compression. Most existing server protocols perform many of the presentation layer functions. An example of the syntax conversion process is the Virtual Terminal Protocol (VTP) which normalizes code for carriage returns, line feeds, etc., between dissimilar terminals, so they may communicate. Perhaps a better example is the File Transfer Protocol which allows the movement of files between heterogeneous computer hosts--clearly a major purpose of networking.

Existing nonproprietary server protocols include the Simple Mail Transfer Protocol (SMTP) defined by MIL-STD 1781, File Transfer Protocol (FTP) defined by draft MIL-STD 1780 and

RFC-765, The Trivial File Transfer Protocol (TFTP) found in RFC-783, Telnet (MIL-STD 1872 and RFC-764), and the Name Server Protocol defined by IEN-116, to name a few. A DoD Network Virtual Terminal (NVT) protocol and a new Name Domain Server are due out soon. The specification for the WIS-CUS calls out the Telnet FTP and a TCP SAP (Service Access Protocol) protocol, leaving the other necessary servers up to the contractor. This is sure to guarantee WIS unique protocols at the most critical layer for interoperability. The WIS Functional Description does call for SMTP, but it is of little value if it is not in the CUS specification. Also, will the planned WIS Joint Mission computers be dependent on the CUS vendor selection/invention of server protocols? While there are probably sufficient data right clauses in the CUS contract, how will the WIS SPO ensure these protocols are machine-independent?

There are proven server specifications widely implemented and in the public domain for most applications. Two that are not very available are complete data base management and file access server protocols. No one is sure how encompassing the former should be, but the latter is being worked by ISO. Where FTPs necessitate the transfer of a whole file, the OSI File Transfer, Access, and Management (FTAM) protocol will allow access to and the transfer of individual records from large files.

Encryption where/if applicable should also be done by the presentation layer as an adjunct to code/syntax conversion. DES encryption for privacy may be appropriate for WIS, considering the operational modes available until true multilevel security is obtained. However, there is a debate on where encryption belongs. Many put it in the transport or data link layers as a matter of convenience. However, effectiveness may be lost if it is this far down the chain of protocol layers. Other user programs may become aware of the encryption and be able to intercept and copy the encryption key.

7. The Application Layer (Layer 7)

The application layer contains the files, applications, and meaningful (useful) information stored or accessible by the computer host. At this time there are no universally specific requirements for application layer protocols.

8. The ISO Reference Model

The ISO reference model is a hierarchical set of protocols with seven independent layers. Many of the protocols cited were actually developed before final definition of the ISO layers, and there may be some overlapping of functions between practice and the reference model. This is especially true of the upper level protocols, levels 5-7. Also standards, procedures, and specifications were referenced when discussing specific protocols. One does not implement standards. While the standard may be applicable to all networked computers and available to all users, the implementing specification may be vendor- and/or machine-unique.

C. DEPARTMENT OF DEFENSE PROTOCOLS

The Department of Defense (DoD) has mandated the use of DoD protocols through OSD memos, letters, and through military standards. Although these DoD standards greatly facilitate the movement of data, they do not yet represent a complete protocol suite. For example, server protocols are still in the approval or development stage. Security protocol applications are being studied by NSA; virtual terminal protocols are being developed and there exist many server protocols which have not yet begun the process for formal approval. This is especially important in the LAN environment where the sharing of special and dedicated resources on the LAN is essentially the reason for having a LAN. A domain name-server protocol is being developed by ISI and should be released in the near future. It is expected that these and other protocols will gradually be added to the DoD

list, and as they are, they should be incorporated into the WIS protocol suite where appropriate.

There are five advantages to complying with the DoD protocol structure:

a. They are standards. They have been sanctioned by the Office of the Secretary of Defense and adopted by all Military Departments. If the WIS community is to have a near real time and open, albeit secure, exchange of information with the Services' C^2, intelligence, logistics, and personnel data bases, it must adopt the same standards. Also, a common base for training--widely available, competitive software and hardware maintenance of off-the-shelf equipment--results in significantly lower costs to the developers and users for standardized products.

b. They exist. Many companies already offer these protocols as off-the-shelf products; many more are to follow. Most mainframe vendors are developing TCP, IP, FTP, etc., in-house or through government contracts. Also, many vendors in the aftermarket or peripheral ADP business have, or are in the process of developing, DoD standard protocols. These include companies such as ACC, Exelan, SDC, Ungermann-Bass, Wollongong, and Internet Systems.

c. They work. The DoD protocol suite has been tested under stress in many user environments, modified accordingly, debugged, well-documented, and represents a mature set of protocols.

d. They are nonproprietary protocols. Not only are they DoD's standard protocols, but they are fast becoming popular with vendors. Systems and Software magazine stated that "...most companies view XNS, TCP/IP, and SNA as the most important protocols today..." in a special 30+ page report on communications standards

featured in their March 1984 edition. Of the protocols listed, only TCP and IP are vendor independent.

e. They are being developed to the ISO open architecture, and have already and will continue to influence the international standards sanctioned by the ISO.

D. IMPLEMENTING DOD PROTOCOLS IN WIN

The five advantages above bring us to the point of asking why WIN has not yet begun to comply with DoD standards. The point has been studied for the past two years, but no strong action has been taken to put the results of these studies into action. Enhancement of a unique, nonstandard protocol has only short-term gains, and it delays and makes more difficult the transition from the NCP-based WIN to the TCP-based WIS. The development and implementation costs for these changes could easily be applied to a TCP/IP solution.

1. Translation From NCP to TCP

It has been suggested that retaining the WIN NCP protocols and translating them to TCP would be more effective than a complete conversion to TCP. The following addresses this argument.

a. Introduction. Protocol translation is an attractive alternative in many network situations, especially where a considerable investment has been made in a prior technology. It solved the problem of compatibility between the old protocol and the new during a transition period. Translation may protect the investment in the old protocol without requiring any hardware or software modifications, which could be difficult and costly.

For WIS, this would mean connecting a processor of some kind to each H-6000 computer. The processor would interpret the NCP protocol stream from the H-6000, provide the proper command and responses to it, and support a TCP data stream over the network. In principle, this would minimize the need to

modify H-6000 communications software, which is an advantage in terms of cost and reliability.

On the other hand, no protocol translation approach is without its problems. The sections below detail several problems in attempting to build a translator between NCP and TCP. Its performance is likely to be poor, its cost high, its functionality limited, and its reliability unacceptable. We recommend against further consideration of this alternative.

b. Poor Performance and High Cost. Consider even the simplest case, that of translation between the two character sets, ASCII and EBCDIC. It is not possible to make a one-to-one translation, because each of the two sets includes special codes not present in the other. Thus, there are some ASCII codes that must be represented by two EBCDIC characters, giving the first of the characters a special "escape" meaning. That may not sound too bad, but consider the implications for a highly interactive user program. Users will frequently have to type twice as many characters when using an EBCDIC terminal and the translator as when they use the ASCII original. This is just a simple illustration of the first drawback of protocol translation, the penalty in performance and operating cost that should be expected. This penalty can be significant.

Another source of performance problems is that the connection between the two H-6000 computers is much more complex than a single NCP connection or a single TCP connection. It is actually three connections: an NCP link from the first H-6000 to its translator, a TCP link from the first translator to the second, and an NCP link from the second translator to the second H-6000. Coordinating these connections is problematic. For example, providing effective flow control across these three connections and among the four elements is difficult. The natural difficulties of coordinating four different processes in four different computers are exacerbated by the differences in the flow control mechanisms in NCP and TCP.

c. Limited Functionality. The functionality problems that can be expected in translating from NCP to TCP are much more severe than those described above for the ASCII to EBCDIC example. Instead of the issue of not having an exact correspondence between character sets, there is the much thornier problem of not having an exact correspondence between network protocol functions. TCP can do some things that NCP cannot do. NCP can do some things that TCP cannot do.

The features common to both protocols include checksums, error detection and correction, opening and closing connections, flow control, fragmentation, and interface with the subnetwork protocols.

There are NCP features not present in TCP. The WWMCCS requirements included operator interface, statistics reporting, WIN data collector interface, accounting procedures, and network job initiation. Strictly speaking, these functions are not part of the NCP protocol, but rather the WIN NCP implementation. The fact remains that any protocol translator must address these features in addition to providing a basic TCP implementation. Critical functions would be lost otherwise, and the acceptance of the resulting system by users would be limited.

On the other hand, TCP/IP provides for functions not present in NCP. These include datagram service, precedence of communication, and internetwork communication. There is just no simple way to provide for these functions in NCP.

There is a WIS NCP feature not present in TCP that would not need to be translated. The message retention facility is a WIN-unique strategy for protecting against message loss. A WIN message is acknowledged as successfully transmitted before it is delivered to the NCP. To protect against message loss by the intervening hardware and software between the IMP and the host, the NCP saves the last "n" messages on disk. This will not be necessary with TCP, because TCP acknowledges at the transport level.

d. Unacceptable Reliability. In any translation between two protocol environments, one of the stickiest problems is that of address translation. The basic difficulty lies in the ground rule that the NCP implementation should not have to be changed, and yet it should be able to address protocol servers and services that it could not address before. Think of this in terms of dialing a telephone call to France. It is not surprising to us that we have to learn an extra set of dialing codes even to reach France, and that we have to have the correct number for the person we want to reach. An automatic translation would have the job of converting from a number which appears to be "local" into the format required by the "remote" protocol. Most solutions to this problem are awkward and error-prone.

Another source of unreliability is the inherent complexity in translating between two network protocols, each with a complex state diagram, with many events and transitions. There are literally hundreds of possible sequences of events to be considered. Many of the unusual event sequences are not likely to be handled correctly in both protocols. If they were designed and programmed correctly, it is likely that they were not thoroughly tested and debugged.

A protocol implementation is a collection of many actions in response to events. Necessarily, there are many special cases and rare situations. Testing a protocol implementation is a demanding job. All of these characteristics are magnified enormously in protocol translation. A full-function protocol translation between NCP and TCP might not be completely debugged until several years of operation have passed, years likely to be fraught with lingering bugs and mysterious problems.

e. Summary. Protocol translation between NCP and TCP is likely to be inadequate, inefficient, and unreliable. In the fifteen years of experience we have with packet-switching networks, there are few, if any, working examples of protocol

translators at the level of a NCP-TCP translation. It is not a good idea, and we strongly recommend consideration of this approach for WIS be dropped once and for all.

E. FRONT-END VERSUS HOST PROTOCOL IMPLEMENTATION

1. The General Case

In implementing network protocols on large, general-purpose computers, the question often arises as to whether the protocol should actually be implemented in the host or elsewhere. The general arguments in favor of a host implementation are:

- The host has a general-purpose operating system, with a scheduler, process handling, large memory address space, and other facilities helpful in implementing a protocol.
- Some host implementations can take advantage of the fact that the user process is in the same computer as the protocol process and thus avoid inefficiencies such as copying data blocks unnecessarily.
- Sometimes a host implementation is the only way to support true host-to-host process-to-process communications. Any protocol implementation outside the host must emulate some other kind of device (e.g., a terminal, a disk, etc.)
- Performance may be improved, and operations are certainly simplified, by staying with the simplicity of a single host implementation of the protocol as opposed to involving another computer in handling the data stream.

On the other hand, the balance of opinion in the network community has been shifting over the years to favor the implementation of many protocol functions in dedicated processors. This, of course, is just one example of an industry-wide trend to specialized processors of all kinds, as the microelectronics revolution has dramatically lowered the costs of processing power.

At the same time, there is a growing awareness in the computer industry, and especially in DoD, that the majority of the life-cycle cost of most systems is in the software, not the hardware, and often comes in the maintenance phase of the system. Thus, planners are tending to trade hardware dollars for software dollars, investing in equipment to lower software expenses. Thus, the arguments for out-board implementation of network protocols in some kind of network front end (NFE) are as follows:

- Programming large hosts is expensive. Reprogramming them is very expensive. The best way to protect this investment is to leave it alone.
- Micro-based NFEs can be 10 to 100 times more efficient in handling network protocols than host computers because they can be optimized for this particular task.
- The programming in the NFE can often take advantage of more modern software production techniques, and can be isolated from the host software production and test environment.
- The software in the NFE can be changed without changing the host software.
- The host software can be changed without changing the NFE software.
- The NFE software can be a generic implementation of the protocol, supporting several different kinds of hosts. With the addition of several unique host-specific interfaces, one implementation can do the work of several. This lowers costs, speeds schedules, and reduces risk.

2. The WIS Case

Should WIS implement TCP/IP in the H-6000? It is certainly not necessary to install TCP in the host; an implementation in an NFE is a practical alternative. In fact, considering the significant value of the software investment in the H-6000s already, it is advisable that the network protocols be implemented elsewhere, merely on the basis of leaving that software investment undisturbed.

The implementation of a major network protocol like TCP into an older computer like the H-6000 can be a very traumatic experience. Older operating systems like GCOS are not designed with networks in mind, and do not have many of the inter-process communications primitives that would simplify the implementation of TCP. GCOS has large numbers of software routines and complex interrupt structures to service many resource requests. A GCOS implementation of TCP would have high risk.

A front-end operating system does not need this complexity or generality. Off-loading TCP to an NFE frees up additional computer resources in the host, and provides highly efficient communications based on large blocks of information.

TCP under GCOS would need to have its own mini-scheduler to decide which process gets to use the network access line next. It would manage state tables, multi-level queues, and memory mapping. It is almost like an operating system within an operating system. Experience shows us that such implementations work best on mini-computers and micros and worst on large mainframes.

Furthermore, while there are no technical advantages to the H-6000 TCP approach, there are definite advantages to the NFE TCP approach. In WIS, it may be necessary to have TCP support for the Level 6 as well as the H-6000, also as well as a variant for the Datanet 8. With TCP in the NFE, the network and higher-level applications could be all the same, with unique host-specific interfaces. There are many practical advantages here.

Implementing TCP in some kind of front end also gives the opportunity to facilitate the WIS architecture. The NFE may have the capability to serve as the interface units on the LAN.

Obviously, in this case, the TCP would be implemented on a modern microprocessor-based interface unit. This has advantages for reliability, maintainability, and modular improvement over time. The choice of the most modern hardware and software base

for TCP will make life much easier during the transition. This is also the best choice in light of the need to upgrade eventually from TCP to TP (see the Section G below).

The only caution that we would urge about off-loading TCP to an NFE is that "there is no such thing as a free lunch." It is not possible to reduce the impact on the H-6000 to nothing. Any TCP implementation in a front end relies on a good protocol between the host and the front end. In fact, the overall network service to users is limited by the least capable of the two protocols, TCP and the host-to-front end protocol. It does not matter if TCP is highly reliable and offers good performance if the host-to-front end protocol does not. The front-end protocol implementation must be as good as the TCP implementation. They should be "performance tuned" together, and tested for functionality and reliability together.

There is a Catch-22 here. The host-to-front-end protocol must be as good as TCP, yet the whole idea of putting TCP in the front end was to off-load that function from the host. This is a challenge to software designers--how to make a front-end protocol that is a worthy match to TCP but uses significantly less computer resources, and has substantially less implementation risk and software maintenance cost. This is the area in which most unsuccessful protocol front ends have failed. This is the area we recommend that WIS study most closely.

F. TRANSITIONING FROM WIN TO WIS

The dissimilarities between the present WIN system and the planned WIS highlight the impact rapid technological advances have made to command and control systems in the last eight years. Unfortunately, this same impact plays its toll on transitioning from the WIN to the WIS, which might be best described as forced evolution.

1. The Requirement

To begin with, there is no merit in continuing with a unique WIN protocol structure. The DDN's roots are in the ARPANET, which for thirteen years had been optimized around the NCP protocol. The DDN is now essentially a TCP/IP-based system and is being reoptimized around its new protocol base. The WIN is presently an NCP-based adjunct that will not benefit from improvements made to the TCP/IP network.

The requirement is to transition the WIN to DoD protocols in order to put the WIN in line with current standards and to make its processors compatible with CUS NADB and JMH processors. The establishment of a standard communications foundation will facilitate the migration of H-6000 programs and the development of new programs in support of JOPES and NISROC requirements.

2. Interfacing the H-6000 to the WIN

As discussed above, translating NCP or implementing DoD protocols in the host are not recommended solutions. The preferred choice is to interface the H-6000 using front-end processors, which would minimize the software changes required in the H-6000. Much work, in the form of tests and studies, has already been done to this end by JDSSC. The solution space has been narrowed to two alternatives. The first is to front-end the H-6000 with a Honeywell DN-8 processor and then interface the DN-8 with an SDC interface unit to provide TCP/IP and LAN or DDN access protocols. The hardware already exists for this solution, but some software modification is required to interface the DN-8 with the SDC interface unit. Development should be complete by the end of 1984 and testing is scheduled for early 1985. The disadvantage of this solution is that it is expensive. Two interface boxes are required, and software maintenance charges for the DN-8 are expected to be high (about $12,000 per year per processor).

A more economical solution with higher technical risk would be to interface the H-6000 with the two boards developed

by Martin Marietta under an RADC contract. One board would be installed as part of the H-6000, the other would go with the SDC interface unit. Lower purchase and maintenance cost are expected with this type of approach. However, Honeywell has stated that this approach uses proprietary software and may or may not allow its use. This problem needs to be resolved if it is still an open issue.

3. Interfacing the Remote Network Processor

Honeywell Level 6 computers serve as remote network processors (RNP) on the WIN and also perform the function of terminal concentrators. The RNPs would also have to be interfaced to the DDN and WIS LAN. Currently, all RNPs are connected directly to their H-6000 host by dedicated circuits. These dedicated circuits should and could be replaced by DDN or WINCS common-user circuits. But there is a problem.

In order for a terminal to access a network, it must first log-on to and access its parent host which will pass the terminal-user's identification number to the remote host. If authorized to access the remote host (prearranged), the connection is completed. This log-on and access procedure is uncontrolled if terminals, via RNPs, have direct access to the network. The problem should be solved with JDSSC software release RNP 3.0. It is recommended that this release be monitored by the WIS JPMO.

The RNP is a level 6, which is very similar to a DN-8 with different software. The SDC interface unit for the DN-8 could also serve as the interface unit for the RNP except for the major software differences in the level 6 RNP.

4. The Protocol "Or" Box

The COINS program is getting ready for a long transitioning from NCP to TCP. Their problems are different from those of the WIN, and it will involve a different NCP and different hosts than those used on the WIN. To ease their transition, they have built and are now testing an interface unit that contains both TCP and NCP. A sending host can select NCP or

TCP depending on the protocol supported by the receiving host. This approach may or may not be applicable to the WIN transition, and has been considered by MITRE. While this work appears to be progressing satisfactorily, this approach would involve more new hardware and interfaces, and would be a major development effort. The NCP in WIS is not as neatly modular as that in COINS.

5. The Transition to WIS Upper Level Protocols

The CUS work station will be able to access H-6000 computers by emulating a Honeywell VIP 7705 terminal. However, this seems to underplay the capability of the CUS LAN architecture. It is therefore assumed that some of the present WWMCCS application programs will be exported to CUS hosts. This could be a nontrivial job. Planning should be concurrent with NCP to TCP planning, and implementation should be tried as soon as the first CUS LANs are in operation.

WIN should be in the process of converting to TCP about the same time WIS LANs are being installed. The WWMCCS hosts will lose some of the functionality they now have with NCP. These functions will have to be reimplemented in the H-6000 hosts in upper-level protocols, if they are to be retained.

After the network is converted to TCP, it will be time to start planning the evolution to TP.

G. TRANSITIONING FROM TCP TO TP

The prospect of eventually evolving from TCP to the International Standard Organization (ISO) TP is an important milestone for WIS. This section discusses some of the considerations in developing a transition plan.

1. Background

First, TCP/IP is now a mature protocol with many installations, and there is a sizable body of practical experience with it. It was developed by a single group of people with strong

central direction by DoD. It is now available "free" on many computers,including many DEC and some IBM models.

The ISO protocol effort began much later than TCP/IP. It is basically a consensus-seeking process, with no central manager. Many members of ISO have conflicting interests, and so TP has been defined with five different classes of service which can be thought of as five separate protocols. Only TP4 is of any interest to WIS and DoD generally. When TP is mentioned here, we refer to TP4. TP4 is the only class which incorporates full handling of misordered packets, as opposed to detecting packets out of order. It also provides checksums, deals with a datagram environment (such as is found on many LANs), and has other features needed in WIS.

There are some minor differences between TP and TCP, but they are similar in the functions they perform. However, it is important to note that TCP is similar to TP, but not the combination of TCP and IP. The IP functions are not present in TP. Differences between TP and TCP include the following:

- TP numbers packets; TCP numbers bytes
- TP does not provide for the sender to "probe" continually to see if a closed window has been opened; TCP does.
- TP has only abrupt and potentially disruptive termination of transport connections; TCP has an orderly and graceful termination.

There is also an ISO proposal for a protocol similar to IP, termed the "connectionless network service." This is a datagram service at level 3. This protocol offers the same functionality as the 24-bit header for datagrams, but does not include other functions commonly lumped together under the heading "IP," such as ICMP and the gateway protocols EGP/GGP. ISO does not appear to be considering the latter protocols at all.

Instead of the DoD standard 32-bit source/destination address field, the ISO proposal has a much longer variable length field, typically 32 octets long, since X.25 network addresses need to be included as components. The X.121 address for such networks is 10 or more digits long. The IEEE 802 standard for LAN addresses calls for a 48-bit address. Since there is no single numbering authority for international networks, the only practical approach for ISO is to adopt a hierarchical numbering scheme corresponding to domains of authority. This implies very long address fields.

2. TP Versus TCP

WIS should evolve to TP. There are several reasons for this recommendation:

1. It will become the international standard. There is no technical reason for WIS not to adopt this international standard. The WIS requirements at this protocol level are similar to those of many networks.
2. TCP/IP is likely to become an evolutionary dead end, as the limited vendor support it now enjoys slowly erodes.
3. Vendor-supplied implementations of TP are likely to become widely available at reasonable prices, since they will be "standard" products, much as X.25 support has become a standard vendor offering.
4. Interoperability with other non-WIS DoD hosts will require TP in the future, as it requires TCP today.
5. Interoperability with European computers, especially in the NATO community, will require TP, since those countries support the ISO standard.

The arguments against using TP usually come down to inertia--it will be too expensive or difficult to switch from NCP or TCP to TP. These are not arguments against switching to TP, they are arguments about when to switch, how long it will take, and how much it will cost. It is not our recommendation to stay

with TCP forever, nor do we recommend that the transition to TP take place immediately. We recommend a conservative schedule.

How long should this transition take? How should it be accomplished? What are the "hidden" pitfalls and difficulties to avoid? Of course, no one can answer these questions today, because no one has ever converted an operational TCP implementation to TP, much less an entire system such as WIS. The next section considers some useful precedents.

3. Historical Precedents

It is always a good idea to try to learn from history, to avoid making the same mistakes people have made before us. The transition from TCP to TP can be compared to the transition in DDN from NCP to TCP. In march 1982, the DDN Program Management Office (PMO) announced that all DDN hosts must convert to TCP by January 1, 1983. Now that this conversion is over, we can look back on the experience, and draw some conclusions for WIS.

It is important to bear in mind that the DDN conversion actually took place in the research environment of the ARPANET. Their experience is not entirely indicative of a protocol cutover in the operational military environment of WIS. Nevertheless, we can make the following observations on the problems with this project:

- The 9 months allocated for the phase-in of TCP was not nearly enough time. Many programs were being modified over the Christmas holidays in time for the January 1 deadline, and many others were not ready in time.
- January 1 is not a good date for a major cutover of this type, since many people in many organizations must be involved, and coordination is made more difficult by the holidays.
- It is essential to have a good means for certifying a new protocol implementation. In the DDN case, there

was not enough attention given to providing users with access to TCP certification tools.

- It took many months to shake out all the bugs, and some hosts were off the network for months as a result. (Of course, any interruption of network service would not be acceptable in WIS, but the fact that other host systems were forced to accept this penalty indicates the magnitude of the difficulties involved.)
- The use of SMTP as the standard mail protocol was mandated at the same time. This only made things worse, since several protocols were changing at once. If possible, WIS should restrict itself to one change at a time.
- No prior study was conducted of the expected impact of the protocol change on network traffic. In fact, the traffic in the ARPANET doubled as a result of this change! Obviously it is essential to understand as fully as possible the consequences of the protocol change before carrying it out.
- Relay hosts were developed to provide for the transfer of mail between NCP hosts and TCP hosts, but they did not function very well. The interoperability provided for mail, and for terminal access was less than complete.
- Finally, and most importantly, the DDN management made the tactical mistake of treating the transition as a problem that users of the network, the staffs at the host sites, should work on. In hindsight, it would have been much better to present them with packaged solutions. After all, it was really a project for the vendors of TCP software, who are not necessarily the same groups as the users. This lesson will be especially important for WIS.

On the other hand, the DDN experience provides us with some positive lessons as well. Several aspects of the project went well, and deserve to be emulated in WIS:

- The Terminal Access Controllers (TACs) were programmed as dual NCP/TCP devices. Users could select either protocol, enabling them to access hosts running either protocol during the transition period.
- NCP was disabled in the TACs after January 1. This made it impossible for many users to continue to use NCP.
- Furthermore, the IMPs were programmed not to accept NCP traffic after January 1. If this had been possible technically, it is quite likely that many of the NCP hosts would not have converted to TCP.
- As a part of the transition plan, a series of test days was held, one a month for a few months preceding January 1, to try out the entire network on an all-TCP basis. The IMPs and TACs were set not to accept NCP traffic. Measurements were carried out. But the primary effect was to alert everyone that the transition was actually going to happen.
- The relay hosts, although limited functionally, were vital in providing a link between TCP and NCP. Note that they did this not by translation between the two protocols, but by providing communications at a higher level (Telnet or Mail), and by having all communications flow through an intermediate host (which would actually be identified in the Telnet or Mail messages.
- A series of seminars was held around the country six months prior to the cutover to explain the technical provisions to the systems staff at all interested host sites. The test days, relay hosts, and other functions were described. This was very useful in raising the people's awareness of the change, and to show them what to do.

In summary, the DDN experience teaches us:

1. The importance of a long and carefully planned transition period (18 months would have been more suitable for DDN, longer may be needed for WIS);

2. The role of the IMPs and the TACs in (1) providing dual access to the network; (2) turning off the old access during tests and after transition;
3. The value of a good certification program for the new protocol;
4. If the transition is a few hosts at a time, there will be a need for relay hosts or their equivalent or complete interoperability between the two protocols;
5. The need to get schedule commitments from the vendors of the protocol software, not just the users.

4. Recommended Approach

We suggest the following as the way to transition WIS from TCP to TP.

1. Begin transition July 1, 1989; end transition June 30, 1991.
2. Review these dates annually to ensure that vendor TP software will be available (without special WIS procurement--a big practical advantage) before the July 1, 1989 start date.
3. Begin now the development of TP certification facilities.
4. Begin now an analysis of the feasibility of a translating gateway between TCP and TP.
5. Leave room in the BIU for TP, to coexist with TCP.
6. Continue as planned with the present transition from NCP to TCP.

In our view, the key technical uncertainty is the feasibility of a translating gateway between TCP and TP. Given the analysis above on the impracticality of a translation from NCP to TCP, why do we recommend a translation here? There are several reasons:

- TCP and TP are much more similar than NCP and TCP.
- With restricted functionality, the TCP/TP translation will be possible. The only question is how much restriction will be necessary.

- Some kind of translating gateway is essential during transition period from TCP to TP, which may last a long time.

H. LOCAL AREA NETWORK ARCHITECTURAL CHOICES

To what type of local area network should the WIN transition? Local area networks may be classified by their topology (ring, star, etc.), transmission medium (broadband or baseband), type of cable (twisted pair, fiber optic, etc.) and method of access (CSMA or token), all of which have different tradeoffs in cost and performance.

1. LAN Topology Choices

There are essentially three types of topologies for LANs: star, ring, and bus.

a. Star networks use a central node to control branches radiating to all the stations in the network. The central node may be a computer system that performs a store-and-forward message switching function. However, telephone technology has recently been adapted to provide a circuit switching scheme, by which the central node established a dedicated path between communicating stations. Star networks have some cost-saving potential, but normally are satisfactory for only low-bandwidth applications. The unacceptable fault with these systems for the WIS is the fact that the control node is a significant single point failure. Another shortcoming is that the central node can do little other work because it is so busy playing operator. A star topology is not appropriate for WIS, because of bandwidth limitations and the central controller that routes all information through the system.

b. Rings are so named because of the circle formed by the transmission system. A repeater must be inserted

in the ring at each station connection point. Token-passing schemes and time-division multiplexing are normally employed with ring networks. A station needing to transmit a packet waits its turn (until the token arrives), then sends its data into the ring. The packet contains source and destination address information as well as data. As the data circulates the ring, the receiving station recognizes its address and copies the data into one of its buffers. The packet continues to circulate until it returns to the source station, providing a form of acknowledgement. Because the ring consists of point-to-point links, almost any transmission medium (twisted pair, fiber optic, etc.) can be used. The biggest advantage of this type network is that it is the most efficient for computer-to-computer or workstation communications as employed by the WIS. Control may be distributed among hosts in the network, with one computer actually monitoring and the others in a backup role. The major disadvantage is that there are active repeater elements in the ring. If one of these failed, or if the ring were otherwise broken, the entire network could become disabled. Both the proposed ANSI and IEEE standards incorporate schemes to recover automatically from single failures. The reconfiguration would occur within a few milliseconds. A less severe problem can be caused when transient errors on the cable destroy or duplicate circulating control tokens. Some small amount of time will be lost until these are sensed by the controlling node, which will issue a new token. Also, because there is a repeater station associated with every station, propagation time increases as stations are added. The advantage of ring topology generally outweighs the disadvantages.

The ring network should be given careful consideration by WIS management.

c. Bus technology uses a single length of Community Antenna Television (CATV) coaxial cable to form a data path. Fiber optical and other point-to-point media do not lend themselves to the multidrop nature of this topology. Coaxial branches from the main stem may also be used to form a "tree" topology, a variation of the bus. Readily available CATV components have made the bus a popular choice for LANs. Bus networks are based on a passive interface and it is extremely easy to tap into the cable. This has the advantage of being able quickly to configure a network and for adding new stations. Each station is responsible for sharing network management, but the failure of a station has no effect on the performance of the network. Coaxial cable is bandwidth-limited, not so much by the cable, but by the methods in which the cable is used. New schemes are being investigated to increase its data rate potential. The biggest disadvantage is that it is not suitable for fiber optical cable. This is not a serious disadvantage now, but could be a significant disadvantage ten years from now as the quest for more throughput, processing power and greater storage and retrieval continues, as we believe it will. Both CSMA and token access methods may be used. As discussed earlier, the token-bus LAN does not appear appropriate for WIS applications.

2. Transmission Medium Choices

Four types of media are in general use for local area networks: twisted pair, coaxial cable, fiber optical cable and free space.

a. Twisted pair is the easiest and cheapest to install, but has poor noise immunity, limited range and very low bandwidth capability. It would only be suitable in the WIS environment for special situations where cost and time might dictate its use as a temporary measure.

b. Free space includes microwave, infrared, laser and FM radio is broadcasted. All four suffer from bandwidth limitations generally with a capacity an order of magnitude lower than coaxial, and it is most susceptible to covert listening. It has the advantage of being used in special situations, normally indoors to connect intra-building local networks, where cable can not be run. Security and low bandwidth appear to make it a generally unacceptable alternative for WIS, except for very special circumstances. One of these special cases may be FM radio in a modile, remote and/or tactical environment. The 19.2 Kbps FM transmission is just the right bandwidth to service a workstation.

c. There are two types of coaxial cable, baseband and broadband. Both have a center conductor surrounded by a dielectric insulator, which is surrounded by a woven copper mesh in a baseband cable and an extruded aluminum jacket in a broadband cable. Both have moderate immunity to noise. A connection to a baseband cable may be made quite simply by piercing the outer mesh and making contact with the inner conductor. This is a real plus for the authorized user, but also makes it just as easy for a surreptitious listener to pick up everything on the transmission media without being easily detected. Connection in a broadband cable is usually accomplished with RF splitters, which are widely used in the television industry. Therefore, adding new new stations is not as easy as with baseband

cable. The cable must be cut, and retuning the network for signal and impedance level must be checked. Broadband cable has an approximate tenfold advantage in signal attenuation with distance and a twofold advantage in bandwidth. Coaxial cable in both its broadband and baseband form could be suitable for WIS application.

d. Fiber optical cable transmission has the most potential, but only recently has it begun to gain acceptability in commercial and industrial uses. Inexpensive LEDs have replaced costly lasers and greatly reduce cost. Optical splitters, used to tap into the fiber optic line, are still expensive, though inexpensive ones are now being developed. Six-way splitters should be available in the next two to three years. The cost of fiber optical cable application is expected to be competitive with coaxial cable in five years. The advantages of fiber optic technology are; (1) very high bandwidth, currently at 500 Mbps and expected to go to 5 Gbps by 1986, (2) very high noise immunity which significantly enhances transmission integrity, and (3) good security, as fiber optical cable is not easy to tap nor does it radiate. Splitter technology is clearly the pacing factor. Fiber optical cable has far greater potential than any of the other transmission media and its only real disadvantages over other transmission media should be resolved in the near future. Fiber optical transmission appears to be the most compatible with WIS requirements. While this cable would be more expensive to install at the present time, it may be cheaper than to install coax and then have to replace it when requirements grow and fiber optical cable costs drop.

3. Baseband Token-Ring Versus Broadband CSMA/CD

Baseband and broadband multiplexing schemes were described earlier, as were token-ring and CSMA/CD access protocols. While there are other combinations available, token passing works best using a ring topology operating at baseband frequencies, and CSMA/CD was designed to work on a bus at broadband frequencies.

The principal difference between baseband and broadband is that the former uses time division multiplex while the latter uses frequency division multiplexing. This gives broadband the disadvantage of not being able to transition to fiber optical cable effectively, but the advantage of having multiple frequency bands or channels on the same cable. Interface units with as many as twenty channels are commercially available. It is possible to send data over one channel, full-motion TV (if 6 MHz channels) over another, and voice on yet another channel.

Making effective use of the bandwidth is another story. In 1981 the IEEE Computer Committee reported that it was reasonable to expect only an effective 40 percent utilization of available bandwidth on a CSMA/CD system. Recently a user reported that "if you set things up so that the normal peak load is 3 to 5 Megabits/sec, then collisions are rare enough so that they can be neglected." On the other hand, a baseband token-ring system works well in heavy traffic, but is slightly slower than a CSMA/CD bus when light loads are on the network. The reason for this is that the station on the token-ring sends a packet and then must wait for the token to traverse the (idle) network before it can send its next packet. The ANSI FDDI standard promises a solution to this problem, but it will be a while before that standard is actually implemented.

4. Costs

Twisted pair installations, which are by far the least costly, will be disregarded because of their unsuitability as a general purpose network in the WIS environment. Fiber optical

cable is currently the most expensive technology, but as mentioned earlier, new splitters should help drive costs down. Another factor that has hindered fiber optics is the lack of standardization. While standards are on the horizon, there are currently no standard connectors, etc., and hence, no quantity savings from mass production.

A recent IEEE 802 cost study showed that baseband and broadband cable installation cost were about the same, about $1250 per user port (CY81 dollars) for a 19.2 to 46 Kbps hookup. Costs for ring installations would be very similar, perhaps slightly higher. When measured by effective use of available bandwidth, the baseband ring cost is more attractive. Again considering bandwidth utilization, but looking five years into the future, fiber optical cable should have a considerable cost advantage.

5. LAN Performance

The maximum throughput of a CSMA/CD bus has already been mentioned as approximately 40 percent in a heavily loaded network. This figure compares to a roughly 95 percent effective bandwidth utilization for a token-ring or token-bus architecture. Under lightly loaded situations, the token-ring and CSMA/CD have about equal delay, with the latter having a slight (1-2 percent) edge. The token-bus fares far worse with about 30 percent efficiency. In a wartime environment the heavily loaded scenario would be the most appropriate, meaning that a CSMA/CD bus system would have to be significantly overengineered and overbuilt to be effective. We do not believe that the "growth to 7 Mbps" stated in the draft WIS-FD-100 represents a significant overbuild. Intersite delay is also a potential problem. Using the figures in the WIS-FD-100 for gateway and in-transit end-to-end delay, WWMCCS transmission will take twice as long with WIS as it does now. In other words, things will get worse, not better with WIS. Intrasite communications and effectiveness should improve with WIS, but to what extent should this take precedence over

getting a long file transfer from an operational command to the National Command Authority?

LAN delays could also have a noticeable effect on character-at-a-time users. Though none seems to be planned for the WIS, the system is now used in the TTY mode and will probably continue to be used in this mode. Round-trip delays for terminal and TTY character-at-a-time users to remote hosts will be unacceptable.

Unfortunately, the gateway delay does not tell the whole story. The long haul network views a gateway as a host and the DDN allows only eight packets in flight between host pairs. In a crisis situation most traffic is focused on one geographical location, say CINCEUR Headquarters. With five to sixteen computer hosts and 100-500 workstations generating traffic out of that location, the one to three gateways servicing the LAN could easily become overloaded. It would be reasonable to assume that much of that traffic would be destined for the National Military Command Center (NMCC) also served by one to three gateways. For simplicity, only one gateway will be assumed at each location. The DDN would consider the EUCOM and NMCC gateways as hosts and only allow eight packets on the network at any one time between this host pair. When the first of these eight packets arrived at its destination gateway, in this case EUCOM, an acknowledgement would be sent to the originating host (the NMCC gateway). When the acknowledgement was received, the EUCOM gateway would send the next (ninth) packet destined for NMCC. However, this "next" packet was the last packet to arrive at the EUCOM gateway from the LAN. All packets being sent by the LAN to the EUCOM gateway between the eighth and this "next" packet would be discarded, as the gateway performs an IP (datagram) function. The point is that a lot of packets could be discarded in the time it took to transmit and acknowledge that first packet. Measurements on the DDN have shown that 30+ percent of the packets

entering a gateway may be discarded due to this phenomenon during the worst hour of a routine day, which does not necessarily coincide with the busy hour. Additional gateways would, of course, produce more favorable results, but the example used would also be much more severe than in a routine DDN day. When the internal LAN delays and transoceanic DDN delays are coupled with the retransmission delays caused by the gateway dropping packet, an intolerable situation may exist.

The above example was used to make two points. First, delays between sites or command headquarters should be given a more thorough analysis with the results compared to the real or perceived requirements. Second, the gateway problem is being resolved...maybe. Resolution of this problem should not be taken for granted. WIS owns the gateways and the DDN looks at a WIS LAN as simply one to three hosts depending on how many gateways are used. The DDN does not know (or care) if there are 1 or 1,000 hosts behind the gateway.

I. OBSERVATIONS/CONCERNS

a. The cable plant for WIS LANs is too tightly specified in the sense that different types of cable and different LAN access protocols can and should be co-mingled based on the requirement.

b. Throughput capacity on the LAN is currently underspecified. It is not believed that a 7 Mbps LAN with an effective throughput of less than 3 Mbps could sustain 5 to 16 hosts and 100 to 500 workstations, plus peripherals during a crisis or mobilization scenario.

c. The WIS is not currently developing successfully an open protocol architecture. ISO/OSI Level-6 protocols must be selected and specified by government in order to have true interoperability.

d. Security planning is not adequate at this stage of WIS development. Evolution must be planned in enough detail to allow growth, change and the injection of new technologies. The to-be-published Security Evolution Master Plan should already have greatly influenced the WIS-CUS and Functional Description.

e. Transportability and mobility are becoming more important in strategic planning. The draft WIS-FD-100 states a need for a future transportable LAN configuration. Given the fundamental nature of the business DoD is in, it is questionable as to why a mobile is not being developed in parallel with the fixed-site capability. Commands like CENTCOM, and probably EUCOM, require a mobile capability first.

f. So far, no significant provisions for teleconferencing, terminal or TTY access appear to have been considered for the WIS. Yet, they all play an important role in the functioning in today's WWMCCS environment.

g. The WIS is required to use the DDN as its long haul communications backbone. Even though DDN usage is a requirement, a detailed plan for integration will require approval. We believe a joint action will be needed to allow the OJCS the opportunity to comment on or modify the planned transition. This can be a long process and could take up to 24 months to prepare the plan and run it through the approval process.

h. The draft Functional Description, WIS-FD-100; the WIS-CUS specification, and security analysis seem to be more technologically oriented than requirements driven.

J. SUMMARY

There are no easy answers to choosing a LAN technology today that will serve the command and control community into the next century. A token-ring baseband network employing fiber optical cable appears to be most promising, but is not as available and is still more costly than current alternatives.

More than one technology should be tried and it is entirely feasible that no single LAN architecture will be optimal for all situations. Only supportability requirements will limit the number of different types of LANs than should or could be installed. Where logistics permit, LANs should be tailored to the user environment and based on performance and cost criteria.

Ten years ago less than a dozen computer networks existed. Today local area networking alone is a $100 million per year business. It is expected to grow to a $1+ billion business in the next four years. This rapid growth will quickly sort out some of the wrong answers. The 200+ networking vendors cannot all survive. De facto standards based on the OSI open architecture will spring from popular vendor offerings and these same vendors will determine the final selection of standards developed by planners in standards organizations, as ANSI, IEEE, CCITT, NBS, etc.

Although the WIS program cannot avoid getting caught up in this sorting of protocols, architectures and technology, good planning, system engineering and management by WIS can alleviate the adverse effects on WIS performance.

The Open System Architecture has come a long way in providing a means for interoperability, reliability, and flexibility through protocol standardization. The protocols that are available today can be expected to change as improvements and new requirements dictate. New protocols will be added to fill existing voids in the seven-layer framework. Complete compatibility between dissimilar computer hosts is at least ten years away.

K. LAN ACCESS STRATEGIES

Our studies of local area networks centered on Applitek's system. The Applitek product line includes interfaces to baseband, broadband, and fiber cable. The Applitek product line

includes the necessary modems and tapes for each of these media. DoD has had considerable experience with broadband systems, and there seems to be some assumption that this media will be used despite the widespread commercial use of baseband. In any case, the Applitek access package is claimed to operate at speeds of up to 10 megabits per second.

Most of the interest in the Applitek product appears to center on its access technique. At present the two most used techniques are CSMA/CD (or CSMA) and token passing. Studies, simulations, and operating experience seem to bear out that the CSMA/CD or CSMA techniques operate well at low channel usage, while token-passing systems work well when the channel is very active. Applitek's access technique is claimed to combine the good characteristics of both of these. It should be pointed out that the Applitek technique is not a mixture of these two, but is a fairly complicated process which involves numbered slots and the assignment of packets from a given user to a subset of these slots (a token-passing-like system can be made by assigning different slots to different users). Users also use CSMA/CD when more than one user is assigned to the same slot.

This system is reasonably complex and a control element (or command device) is used to establish the slot times and make assignments of users to slots; also the system is instituted so that any interface can assume command, and there is a numbering system to determine who takes control when a situation arises where control might be taken by someone.

Applitek has done considerable testing of the system and has verified estimates of system performance. The system is more complicated than most of those commonly used, and this of course raises questions of performance in complicated operating situations, price, reliability, and security.

The performance issues are perhaps more readily dealt with. There are several features having to do with priorities of users, assignment criteria for slots, etc. The problems here

are similar to those in setting up operating systems, and the system performance in a given situation will reflect the tuning of the system. As testing and operating experience increases, any problems here should be alleviated and system performance should benefit from the complexity of the system. (Since the "fallback" can be CSMA/CD or token passing, the system must perform at least as well as those in a given situation, and any problems appear to lie in reaching the right strategy at a given time frame.)

The reliability would concern the extra electronics, and this is a real consideration; however, electronics is becoming more and more reliable, and the system has several features to alleviate failures of access systems. Reliability of the software is another matter, and considerable testing and usage will be required to assure reliable operation in crisis situations.

Security is, of course, another matter. The DoD protocols are now being studied for use with the system, and final resolution of security considerations would have to include studies of the system with these protocols in place.

The Applitek system's ability to reconfigure, if cable or the present controller fails, involves the remaining devices "counting" until some device takes charge of each operable section of the system. This means control can pass around, and the controlling device might be chosen by a mischievous cable cutter or by selectively removing other devices. This would lead to problems in that a malicious user might cause a particular device to take control of a section of the cable. The complexity of the programs would also have to be considered in regard to system verification.

In summary, the more complicated access strategy used by Applitek appears to offer superior operating performance vis-a-vis other systems now in widespread use with regard to cable usage; however, systems must be tuned, security might be a problem, and DoD protocols must be incorporated. Because of

the continuing drop in electronics prices and the continuing increase in electronics reliability, the long-range prospects appear reasonable.

III. WIS LAN SECURITY ISSUES

A. OVERVIEW

The WWMCCS Information System Modernization Program is employing advanced techniques in local area network (LAN) technology to provide a framework for the continuous evolution of information handling facilities at the major U.S. command centers around the world. This LAN will link the existing Honeywell 6000 mainframe computer, the WWMCCS Intercomputer Network (WIN) and specific new functional modules and work stations. Over time, many of the functions presently being performed on the H6000 computers will be transferred to new components of the framework, with eventual full integration of these functions and replacement of the Honeywell systems.

A major modernization program such as this is often considered highly risky when undertaken in the private sector or other portions of the government but the one factor which frequently does not come to play in those efforts that even further complicates the WIS activity is security. Some of the information being processed on the computer systems that support WWMCCS is among the most sensitive, timely, and crucial U.S. military information. In addition to this highly sensitive and perishable information, there are vast quantities of routine logistics and administrative information which is vital to the successful execution of command and control operations. The system is far flung around the world, with thousands of users interacting continuously to keep the various data bases up-to-date. There is no way that all of the people that deal with this information can be cleared to access all of the information contained in the system. But somehow a fully integrated

system providing timely and accurate force projections based on vast quantities of routine and highly sensitive information must be available on a continuous basis.

Much has been written about the technologies being employed in the WIS Modernization Program. Considerable information has been presented over the last several years concerning security as it applies to local area networks from the perspective of end-to-end encryption and trusted network interface unit activities (Refs. 1, 2, 3). There is, however, much more to security in a system like WIS than encryption on a local area network. This paper will discuss several higher level security issues that face the WIS, analyze alternative solutions available today or in the near future, and either recommend particular solutions or indicate early research and development activities required to obtain improved alternatives.

B. BACKGROUND: THE PRESENT WWMCCS ENVIRONMENT--A SECURITY PERSPECTIVE

At the present time the WWMCCS ADP system consists of twenty-eight sites with thirty-five systems consisting of approximately eighty-five Honeywell 6000 processors. There are approximately 76 Level-6 Remote Network Processors serving as terminal concentrators and hundreds of Honeywell and other terminals. Most of these systems were installed in the early to mid-1970s using the then state-of-the-art Honeywell 6000 hardware and the GCOS operating system. Despite much publicity to the contrary, the system has served the functions for which it was intended very well.

One of the significant areas that has troubled this and many other such systems for years is computer security. Physical, administrative, procedural and communications security techniques are well understood and effectively applied in the WWMCCS system. The major difficulty has been and continues to be in the area of hardware and software integrity, the

traditional computer security problem. At present most of the WWMCCS computers are located in Top Secret System High facilities; there are a few systems that operate above Top Secret and there are several at the Secret level with at least one system at the unclassified level.

In the mid- to late-1970s when it became evident that networking the WWMCCS systems together would provide significantly improved timeliness and functionality, the lack of hardware and software integrity had a serious limiting effect on the utility of networking. The WWMCCS Intercomputer Network was created to link the Top Secret WWMCCS hosts in a system high network. But there were no means to include the Secret or above Top Secret hosts on the network. Furthermore, the users of the Top Secret hosts on the network all had to be cleared to have access to any information on any host on the network. The lack of substantial hardware and software integrity measures in the H6000s precluded relying on those hosts to isolate information other than on a very rudimentary basis. Thus, even though the computer systems contained 85 percent SECRET information they were forced to upgrade to TOP SECRET in order to be part of the network.

The vulnerabilities and inefficiencies of this form of operation were well understood in the late 1970s, and solutions to these problems were specified in the requirements for the upcoming WIS Modernization Program. The general approach of the evolution of the WIS system from the existing Honeywell mainframes to fully integrated local area networks supporting modular components performing specific functions is shown in Fig. III-1. The concept of installing a local area network with work stations and specific functional modules at each WWMCCS site attached to the Honeywell 6000 system and the gradual evolution of functions to specific modules on the LAN is probably the only technologically and politically feasible solution to modernizing the WWMCCS system. There is no practical

way to install a turnkey mainframe replacement for the existing systems. The ability of the LAN to provide a framework for the continuing evolution of the system will help preclude returning to the single mainframe limitations that have plagued WWMCCS for the last decade. While this solution is technologically sound, it does pose a number of additional constraints on the establishment of reasonable security solutions for the WWMCCS system.

The LAN framework provides a distributed solution which has many advantages from a reliability and modularity perspective. However, it adds several more components to an otherwise difficult security situation. If, for example, it had been possible to substantially improve the hardware and software integrity of a mainframe computer system to replace the Honeywell 6000s, then that solution would go a long way toward relieving the computer security bottleneck that has existed on those systems for years. Now such a host computer solution must be integrated with a LAN to ensure that information sent from the host either to a local work station on the LAN or to a remote work station across the WIN cannot be modified or compromised in any way.

Among the major issues confronting WWMCCS in this area are provision of secure communications paths across the LANs and WIN. Section C of this chapter provides a review of the security vulnerabilities of LAN environments and the various techniques for overcoming those vulnerabilities ranging from separate LANs for each security level to End-to-End Encryption (E3) and trusted network interface units. A second issue is the question of network security policy. Depending upon the configuration of trusted and untrusted computers at each WIS site, there will be security constraints placed on the LAN to enforce various forms of network security policy. What are the trusted components of the network and what should be the security policy that those trusted components are required to enforce. In light of the fact that E3 and trusted LAN devices

are not readily available, are there ways to minimize the impact of using trusted and untrusted computers on LANs which do not require special provisions in the LAN? Section III-D reviews the security requirements in this area and provides some interesting possibilities.

A third major area of concern, which builds on the previous two, is access control in a system such as this. Individual users today log onto individual Honeywell 6000 mainframes. If they wish to communicate to other computers on the network, they establish connections across the network and are authenticated in mutually acceptable arrangements between common hosts. Extensive procedures for user authentication have been developed over the years in the present system. Equivalent procedures for handling much more complex authentication arrangements will be needed for the WWMCCS local area nets and more efficient and effective ways of passing authentication information to the many different host computer modules on the LANs will be required. Section III-E provides an overview and discussion of this area.

C. OVERVIEW OF LAN SECURITY VULNERABILITIES

The vulnerabilities of local area networks are best described within the overall context of the total systems in which they reside. Perhaps the best description of computer system vulnerabilities is contained in the Report of the Defense Science Board Task Force on Computer Security published in February, 1970 (Ref. 4). This report begins with an excellent description of the nature of the computer security problem, including the types of computer systems and the threats to each of them. Such vulnerabilities as accidental disclosure, deliberate penetration, active infiltration and passive subversion are discussed. In the areas of security protection the report describes typical physical protection measures, hardware and

software leakage points, and communications and organizational leakage points. Figure III-1 is a reprint of the famous computer network vulnerability chart which appeared in this report. This figure describes the various types of security failures that can occur in a computer communications system.

When considering the vulnerabilities of local area networks in the context of Fig. III-1, one can view the LAN as replacing the set of wires from the switching center to the individual users, and/or replacing the communication lines between the processor and the switching center or replacing the entire switching center itself. The LAN is therefore subject to wiretapping, radiation, cross talk, hardware and software failures as well as personnel vulnerabilities from maintenance people, operators and system programmers.

Most of the WWMCCS sites operate at the Top Secret classification level and all users presently having access to the H6000 machines must be cleared to that level, even though close to 85 percent of the information on those computers is classified only at the Secret level. Therefore, a major factor in the evolution of the WIS program is the establishment of a capability to operate at several different classification levels and eventually to allow the use of multilevel secure computer systems.

Since the LAN provides the vital link between all user work stations and computational resources, it is essential that the LAN be able to protect classified information passing through it. This could be done by providing different local area net facilities for each classification or sensitivity level within a WWMCCS site. However, such a solution would be expensive, administratively burdensome and unacceptable operationally. It would be advantageous if the LAN supplied with the WIS were able simultaneously to handle sensitive information at several different levels. Some recent developments in LAN technology are leading to at least limited service of this type with evolutionary expansion possible.

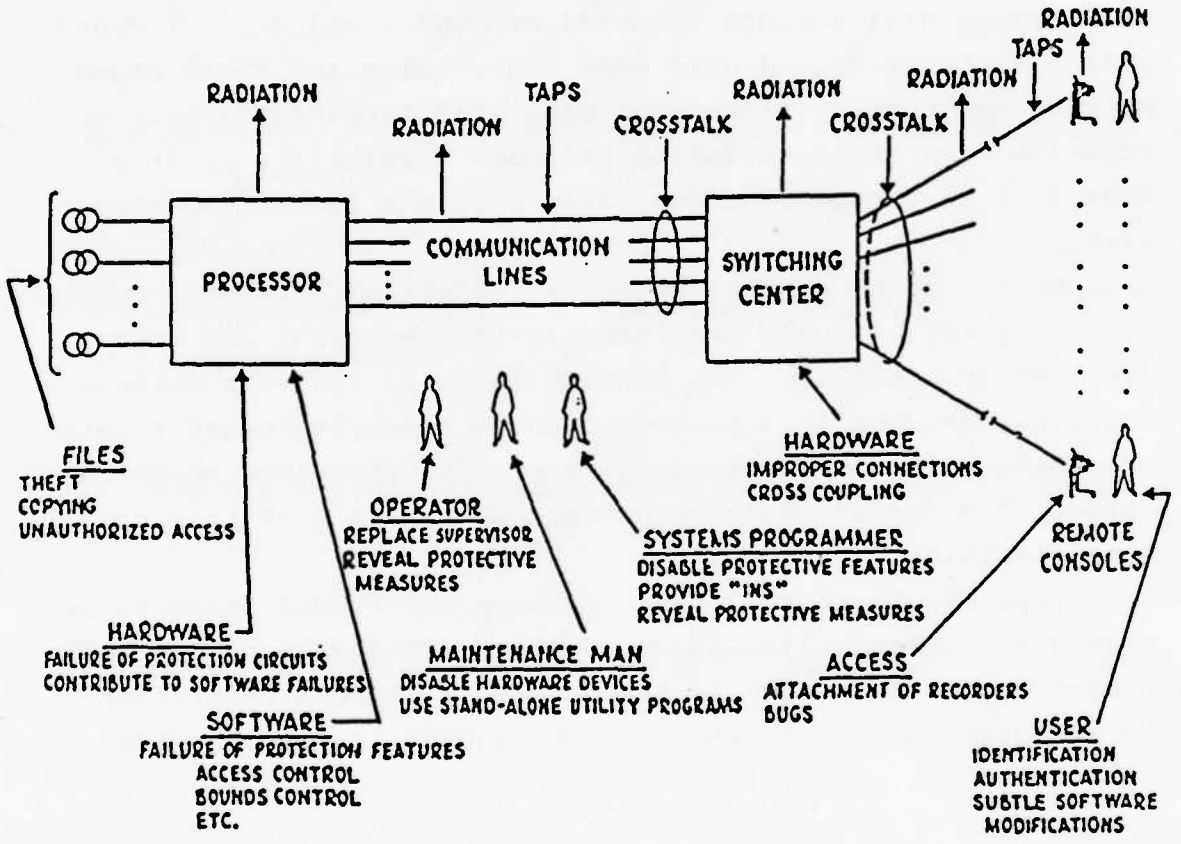

FIGURE III-1. Computer Network Vulnerabilities
(Source: Ref. 4)

1. Security Relevant Characteristics of LANs

A local area network connects communicating devices that are not far from each other but where there is a definite physical separation ranging from a few dozen meters to a few kilometers. In most cases a single coaxial cable replaces an entire network of conventional wires and permits simultaneous two-way communications among communicating devices. Frequently this system will support hundreds of high-speed digital channels with text, data and in some cases voice and image signals. Many of the characteristics of LANs that make them highly attractive from an installation ease and flexibility point of view have very negative connotations from a security perspective.

In a typical coaxial broadband local area network, radio frequency (RF) signals are transmitted over trunk and feeder lines to drop cables which provide links to the user outlets. The trunk is usually a protected cable employing rigid aluminum shielding with a bending radius ten times the diameter of the cable. The feeder cables and drop cables are much smaller and more flexible.

Taps can be placed anywhere along the feeder cable to provide drop cable connections to outlets. The drop cables and feeder cables typically incorporate foil and braid shielding to minimize radiation leakage among cables in close proximity to each other. Directional couplers divide or combine inputs and outputs of RF signals. They are used to ensure that signals being transmitted from any network device will be routed only to the head end device. One of the key advantages of using coaxial cable in systems like this is its ability to support multiple taps. With the isolation provided, each outlet stands alone, and the connection or disconnection of a user device has no effect on the operation of the overall system. Broadband coaxial systems employ components that are readily available from community antenna TV (CATV) or cable TV

systems. These systems are characterized by their ease of connecting new terminal ports simply by adding a tap. It is also easy to move devices anywhere on the network since all devices can view all of the information transmitted over the network.

As can be seen by the above discussion, most of these features provide advantages for the easy installation and operation of a system. But many of them, in particular the ease of attaching taps in an undetected manner and the fact that all information passes every point on the network constitute serious concerns from a security point of view.

Fiber optics is frequently proposed in lieu of coaxial cable in LANs. The characteristics of fiber optics, which are its principle advantages, include higher bandwidth, immunity to electrical interference, lack of electromagnetic radiation, and smaller physical diameter characteristics. However, fiber optic systems are more difficult to tap and amplifiers and related equipment are frequently much more expensive.

2. Protection Measures

This section will describe a series of protection measures that are available for local area network security. It includes an analyses of the overall security of the information system of which the LAN is a part. In this context the physical, administrative and personnel security measures for the LAN are closely linked to those procedures associated with individual terminals and host computers which comprise the total system. The local area network must receive the same protection measures afforded to the highest level of sensitive information contained in the overall system.

In the case of WWMCCS, the highest level of classified information is Top Secret and all terminal, host and communications functions, including the LAN, must be protected to that level. This implies that all personnel must have Top Secret clearances, that all terminals and hosts must be located in facilities cleared for Top Secret, and all communications must either be

contained within Top Secret facilities or be protected via military approved cryptographic devices for handling Top Secret information. Once the WIS system evolves to controlled mode with some users only cleared to Secret, the terminal facilities for those users need only be cleared for Secret information. But, to the extent that the communications linking them may also contain Top Secret information (i.e., on a local area network), those communication links must continue to be protected to the Top Secret level.

There are, in general, two ways to protect a local area network system for handling classified information. The first is to contain the transmission media physically (e.g., coaxial or fiber optic cable) in a protected wire line environment. The second approach is to provide some form of encryption for the communications passing over these LAN media. Each of these will be considered in detail now.

Protected wire line systems typically are expensive to install and they have limited flexibility for attaching additional feeder circuits and drop lines. These systems range from simple welded conduit for use with information with a low level of sensitivity to sealed, pressurized enclosures which have automatic pressure loss detection devices to alarm in case the conduit is broken, accidentally or maliciously, to systems where the conduit must remain under visual scrutiny of security officers at all times.

The second alternative for protecting classified information on an LAN is to provide some form of encryption to the information. Once encrypted, the information is considered unclassified by virtue of having been subjected to the encryption algorithm. Anyone intercepting the encrypted information would have great difficulty in obtaining the sensitive information contained therein.

The simplest form of encryption normally provided to unprotected communication channels is link encryption, in which

the data is encrypted as it enters the transmission medium (i.e., just prior to the modem device) and is decrypted upon being received from the transmission media (i.e., just after the modem device). Unfortunately, due to the high bandwidth of all LAN media, link encryption is not practical. In the case of baseband systems, data rates of 2 to 10 Mbps are beyond the capability of present encryption devices, except for very expensive equipment. Broadband systems which use even higher analog signals are impractical to encrypt. Given that a pair of link encryption devices would be needed at each outlet of the LAN, the cost of such a service quickly become excessive, even if equipment were available today.

A second form of encryption protection for a LAN is the use of end-to-end encryption. In this case the information to be transmitted is encrypted prior to being entrusted to the network and remains encrypted throughout its transmission, being decrypted only upon exiting the LAN interface device. This is a more practical form of protection both from a data rate and cost point of view. The data rates are those of the individual communicating devices and not the aggregate rate of the transmission media itself. In this case the number of devices needed corresponds to the number of drops or communicating devices on the LAN.

At present there are two forms of E3 devices under development for use on wide area networks such as the Defense Data Network (DDN). Once generally available, these devices should be applicable to LANs, though as will be seen, their costs may preclude their use except in very special cases.

The simplest form of E3, the Internet Private Line Interface (IPLI), is illustrated in Fig. III-2. The IPLI and its associated cryptographic device are positioned between the host computer and the network. Groups of IPLIs form communities of interest in which all cryptographic devices share a common key.

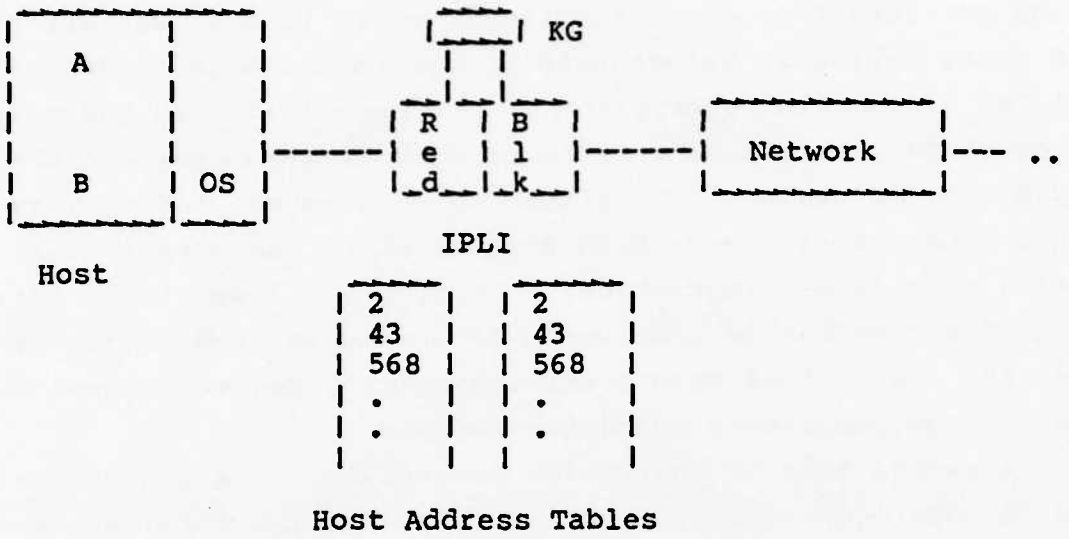

FIGURE III-2. Internet Private Line Interface (IPLI)

IPLIs (and most other proposed E3 devices) come in two sections, a Red side which is connected to the host computer and to the cryptographic device, and a Black side, connected to the network and to the crypto. When the host requests a connection across the network, the Red side first looks up the destination address in its host table to be sure it is a valid member of the host's community of interest. If it is, the Red side passes the data portion of the message to the KG and sends the host table index number for the destination host to the Black side through a special low bandwidth channel which bypasses the crypto device. The Black side constructs a new message header using the index entry in its host table. When the encrypted data is received from the KG, the Black side assembles a message consisting of the new header and encrypted text and sends it to the network. Receiving data from the network works in

the reverse manner. It should be noted that the data portion of the message is protected by encryption at all times except during the brief period the data is in the Red side of the IPLI. The address portion of the original message remains in the clear within the network switches so that they can perform proper message routing.

The IPLI provides a means of isolating communities of interest operating at a specific sensitivity level on a network that may itself operate at a much lower level (even unclassified). Untrusted hosts connected via commonly keyed IPLIs must operate at a specific system high security level. Trusted hosts may operate in more complex and interesting environments as will be discussed later.

3. E3 with Remote Key Distribution

The IPLI provides a valuable capability for connecting communities to a common network. Its disadvantages as presently constructed include the requirement that the communities of interest served by statically assigned (they do not change very often), and the cryptographic keys must be distributed manually and loaded at each site. The hosts in a particular community are not allowed to communicate with anyone outside that community. All communication within the community is protected with the same key so there is no additional protection, over that of the individual hosts in the network, of data among individual members of the community. In effect, each community has its own virtual system high network.

To overcome these drawbacks, efforts to build E3 systems with remote key distribution techniques which would allow host to host, process to process or per connection individualized keying have been underway for some time. With these techniques, a separate Access Controller and Key Distribution Center (with redundant backup) is attached to the network. The fundamental structure of this system is shown in Fig. III-3. The E3 boxes shown are capable of holding a large number of separate keys for use in either a host pair, process pair or per connection basis.

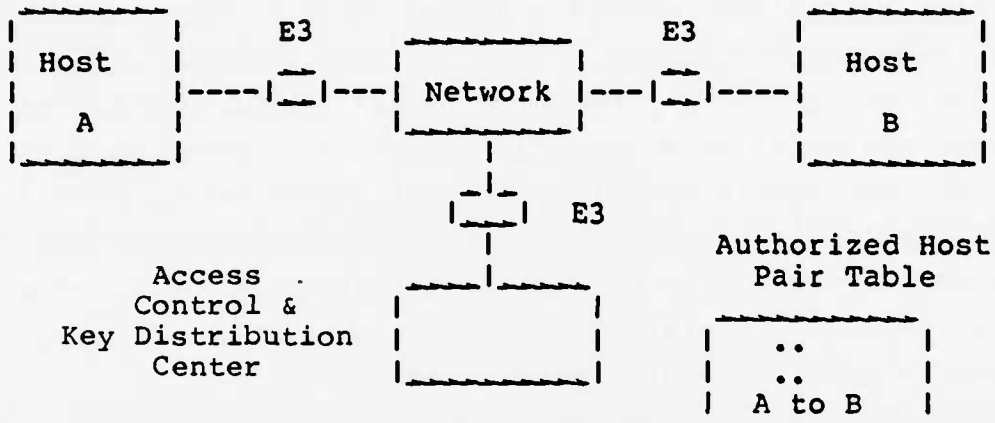

FIGURE III-3. Simple Form of E3 with Remote Key Distribution

The host pair case will be examined first. In this case, when Host A attempts to establish a connection with Host B, the first E3 box checks to see if the key for such a host-to-host connection already exists. If it does, the connection proceeds as in the IPLI case. If the key does not exist, the E3 box establishes a connection with the Access Controller, identifying the source and destination hosts for the requested connection. The Access Controller mediates a decision based on security relevant authorization data in its Authorized Host Pair Table. If the connection is authorized, the Key Distribution Center generates a new unique key for use by these hosts in their communications. This key is passed to both hosts' E3 boxes encrypted in their individual master keys. Once this host pair key is in place, the communication between hosts can proceed as before.

This version of E3 protection has many advantages over the IPLI in that it can operate dynamically with new host pair authorizations enabled merely by making or changing an entry in the Authorized Host Pair Table. This process works equally well with either trusted or untrusted hosts. Untrusted hosts which operate at the same system high security level will have appropriate entries in the Authorized Host Pair Table; untrusted hosts that operate at different security levels or are not authorized to communicate will be prohibited from communicating by the E3 mechanisms. Trusted hosts that are authorized to communicate will have entries in the Authorized Host Pair Table and will control the security levels of their own individual links using their own access control mechanisms as described shortly.

4. More Complex E3 Mechanisms

The host pair connectivity described above provides dynamic authorization of communities of interest and is consistent with the trusted system access control mechanisms described earlier. Nevertheless, at various times the perceived need for more sophisticated E3 mechanisms has been expressed. In these proposals, either process by process or per connection access control is desired. However, when Host A attempts to establish a connection with Host B, it must identify to the access controller the source and destination processes involved. The access checking mechanism becomes much more complex since the Authorized Host Pair Table now must be an Authorized Process Pair Table. All of the update and synchronization problems associated with a host attempting to know the instantaneous status of a remote process are inherent in this type of access checking. These and other factors associated with more complex E3 mechanisms will be examined in detail later.

The problems with the E3 approach are: first, there are no devices available with this capability. The Blacker program is attempting to provide this type of service on a wide area

network such as the Defense Data Network (DDN). Assuming it can be made to work there, it is reasonable to assume that one could extend this capability out to an LAN connected to the DDN. However, initial Blacker devices for the DDN are not expected until 1988 and probably will not be available in production quantities until the early 1990s. The present contract does not call for development of LAN devices although such devices are anticipated to be needed shortly after the initial DDN devices.

End-to-end encryption provides protection for the data being transmitted, but address information, indicating where that data is to be sent on the network, must be kept in the clear. This dual function for the encryption device makes it much more complex than a simple link encryption facility. In addition, some form of key distribution center and access control mechanism is required for anything other than the simplest form of end-to-end encryption. Unfortunately, the procedures for doing key distribution are not fully understood. They are being developed as part of the Blacker program but will not be available in LAN environments until the early 1990s.

While encryption, and in particular end-to-end encryption, shows real promise for providing excellent protection for sensitive information on an LAN, the difficulties to be overcome prior to availability coupled with the cost of the devices themselves, including the key distribution and access control mechanisms, indicated that it will be eight to ten years before this type of protection will be available generally. Therefore, anyone considering LAN service before that will have to provide protection via some form of protected wire line facility.

5. Trusted Local Area Network Capabilities

Because of the problems inherent in applying encryption techniques to LANs and the apparent delay in having such resources available for the next five to eight years, system designers have turned to other means of providing security.

One approach which is straightforward but very limited in its operational characteristics is to install a separate physical LAN for each level of sensitive information. While this approach is immediately available using standard components, it has the severe disadvantages of requiring considered duplication of equipment, extensive additional physical construction of protected wire line facilities, and complex operating procedures to ensure that the user employs the proper terminal and LAN for the appropriate level of security. In systems that involve a reasonable range of sensitive information, this operational complexity may pose such a large security risk that even this low technology approach may be unacceptable.

Several years ago consideration began to be given to finding means to assure that at least the addressing portions of the LAN could be made highly reliable. If the LAN must be enclosed in a physically secured, protected wire line conduit anyway, and if some means could be found for ensuring that the LAN would deliver the information to the proper destination without modification, then simultaneous operation of an LAN with multiple levels of sensitive information could be achieved. The term Trusted LAN was applied to this concept and several studies were initiated to evaluate its feasibility (Refs. 1 and 3).

The simplest form of this trusted LAN is shown in Fig. III-4. In this case the LAN is considered multilevel secure, that is, capable of interfacing terminals and host computers that operate at multiple levels of classification. The LAN interface unit is "trusted" to properly label information entering the network as to its sensitivity and to ensure that information is delivered from the network to destinations at the same sensitivity level.

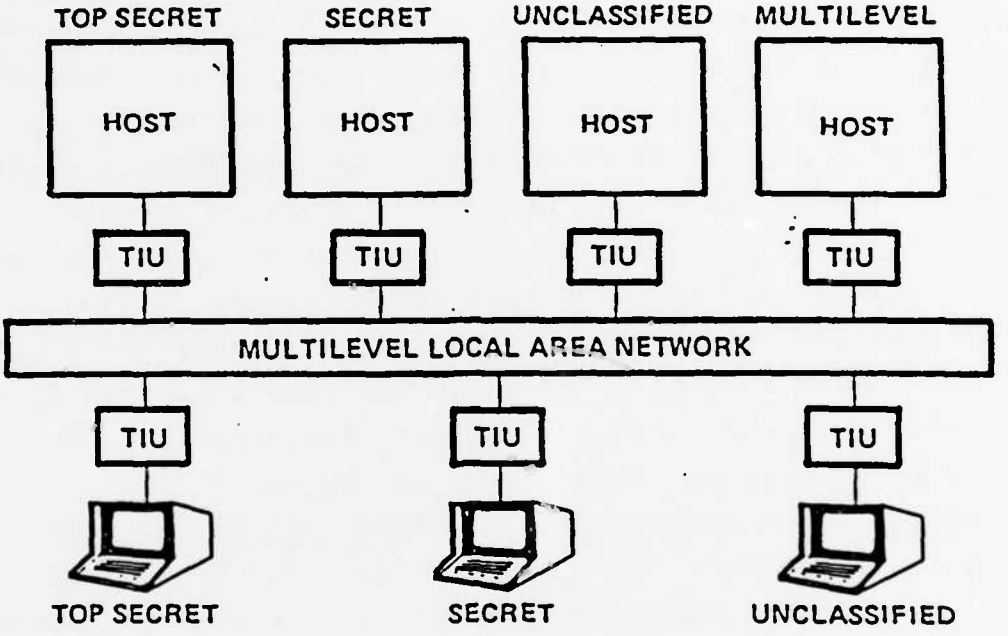

FIGURE III-4. Simple Multilevel Local Area Network
(Source: Ref. 3)

Early analyses of this approach concluded that a trusted interface unit would have all of the development and verification complexity of earlier attempts to build trusted general purpose operating systems. However, more detailed examination shows that much of the complexity that is inherent in trusting such operating systems comes from the need to enforce rather complex security policies that allow users at certain levels in the security hierarchy to read and write data at their present level, to read but not write data at levels below their present level and to (theoretically at least) write but not read data above their present level. This security policy was first stated mathematically by Bell and LaPadula in 1974 and remains the basis for policy enforcement in operating systems today.

However, examining the function which is being performed by the LAN in these scenarios shows that the LAN does not need to enforce a policy as complex as Bell and LaPadula. Rather, the LAN should be viewed as a logical replacement for a bundle of wires directly linking every terminal to every host computer in a local area. Under this view, the trusted LAN must be shown to enforce the much simpler policy stated above of properly labeling information on entry to the network and ensuring delivery only to destinations approved for handling similiar sensitive information.

It is interesting to note that the fear of the complexity of trusted LANs that drove many to favor the use of end-to-end encryption techniques is seriously misleading since in order to properly handle the bypassing of address information in the end-to-end device, a degree of trust equal to or exceeding that of the trusted LAN is required. It is just this issue that is at the core of the difficulty being faced by the Blacker program.

Figure III-5 is a more complex situation involving several physical environments at different security levels with both trusted and untrusted interface units, encryption between the physically protected LANs and host computers operating at all levels of classified information.

In a report entitled, "Cable Bus Applications in Command Centers--Security Issues" (Ref. 2), Robert Shirey describes the WWMCCS LAN requirements and details the use of encryption in considerable detail including an analysis of the potential use of public key cryptography. This report concludes that "multilevel secure LANs can and should be built using Blacker end-to-end encryption," while cautioning that unless action is taken soon (the report date is February 1982), "it is almost certain that the cryptors will not be available when needed." The report acknowledges that local area networks will need both encryption and trusted software.

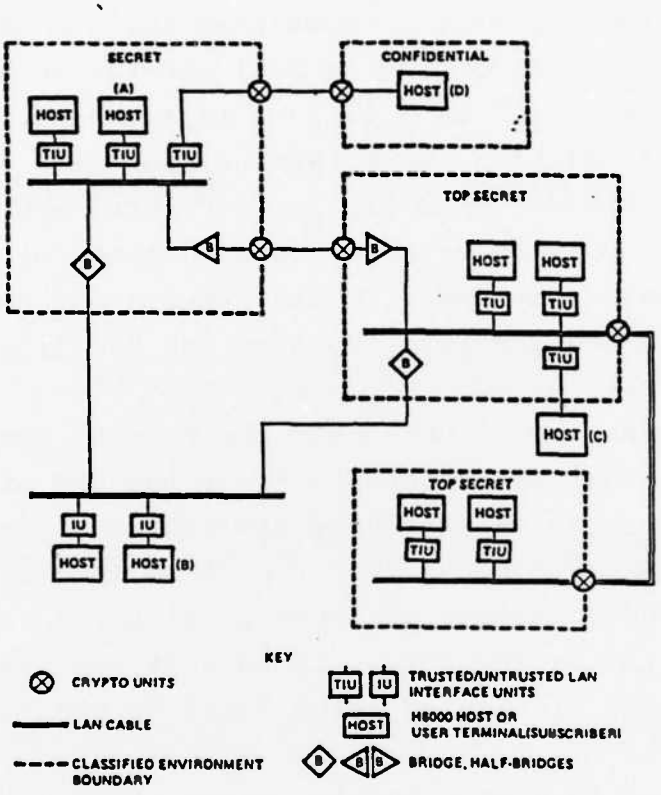

FIGURE III-5. Full Multilevel Local Area Network
(Source: Ref. 3)

A second report on this same area, "WIS Local Area Network Issues," by William H. Blankertz and David A. Gomberg (Ref. 1), describes the LAN security issue in the following terms:

> "By far the greatest risk in the development of WIS LAN is the requirement that the LAN provide multi-level security. It is the expressed opinion of many WWMCCS users that the WIS LAN will be of little utility without MLS. Yet there is no assurance that MLS is achievable; there are no clear-cut solutions to the technical problems that MLS poses. A successful approach to MLS will most likely require features from ... encryption, trusted software and physical protection."

> "The LAN security architecture intersects with almost all aspects of the WIS design. ... Problems such as cryptor availability and lack of adequate tools to develop and verify trusted software may cause considerable delay in the MLS LAN implementation."

An excellent description of the issues surrounding trusted LANs is given in "Design for a Multilevel Secure Local Area Network" by Deepinder P. Sidhu and Morrie Gasser (Ref. 3). Written again from the context off the WWMCCS environment, this report describes the multilevel secure LAN problem, and details the characteristics of several alternative solutions. The report then proposes several operational scenarios for the use of trusted interface units, bridges, gateways and guard systems. The report proposes a series of incremental upgrades, starting from a single subnetwork operating in a system high physically protected environment, and expanding to trusted interface units able to support variable-level terminals and controlled and multilevel mode hosts.

Figure III-6, from the above report, shows a logical structure for a trusted interface unit. The LAN medium, interface, Carrier Sense Multiple Access with Collision Detection (CSMA/CD) and security processor components are all in the physically protected areas denoted as RED. These components all must be trusted to operate correctly. Assuming that this can be done and that these components can be isolated from modification by other portions of the interface unit, then the other components do not have security relevance and therefore need not be trusted.

6. Recent Developments

There are a number of efforts underway to attempt to resolve the local area network security issues. One of the most advanced efforts is the System Development Corporation (SDC) MIL/INT product line which was announced during the AFCEA convention in June, 1984. This local area network system (Fig. III-7) is a broadband cable with a 2 Mbps signalling rate. The product line consists of eight devices ranging from host

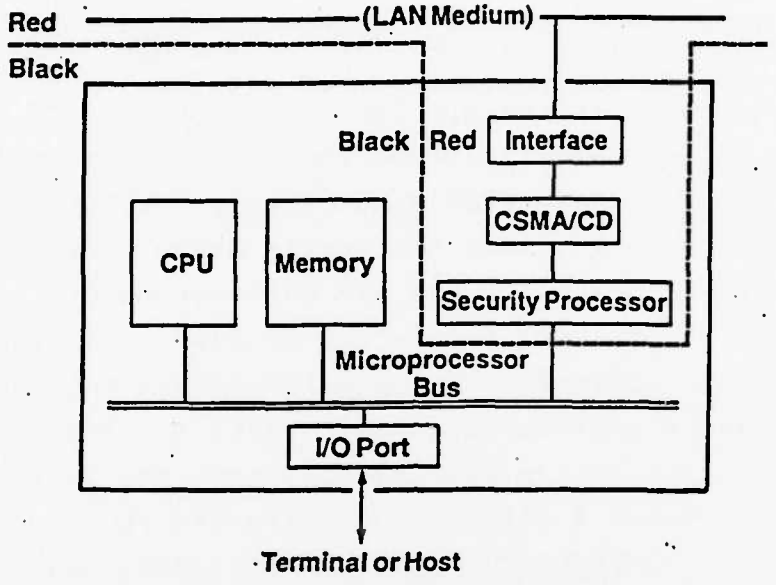

FIGURE III-6. Broadband Secure Multilevel LAN
(Source: Ref. 3)

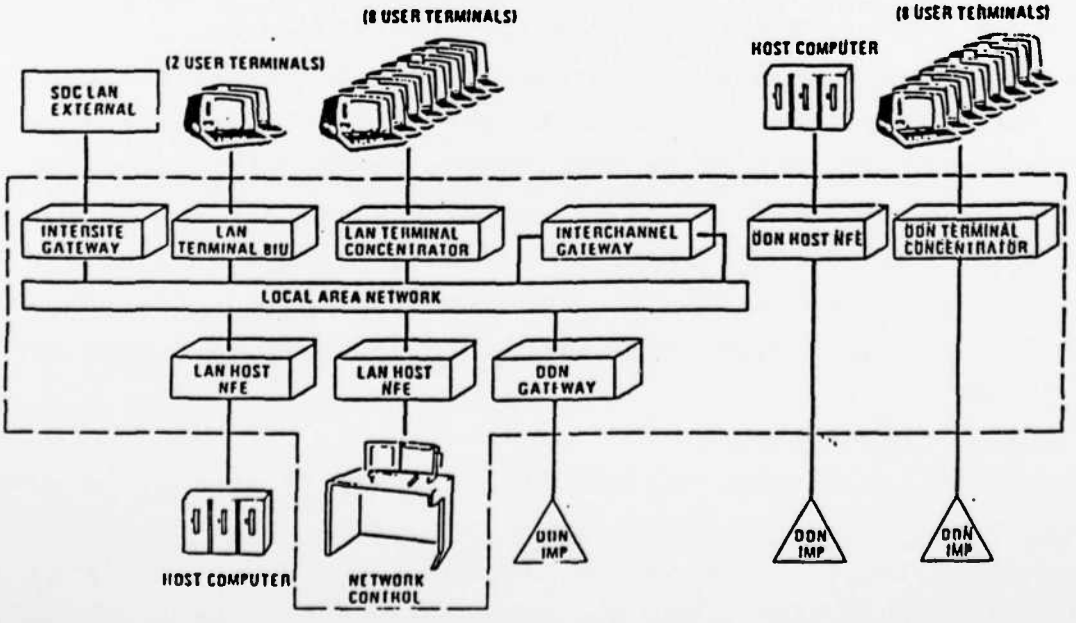

FIGURE III-7. SDC MIL/INT Network Product Line Overview
(Source: System Development Corporation)

front ends and terminal concentrators (for both the LAN and the Defense Data Network (DDN)) as well as intersite, interchannel and DDN gateways and a network control center.

MIL/INT devices are designed to work specifically with the DDN and therefore implement the full set of DoD Standard Protocols including the Transmission Control Protocol/Internet Protocol (TCP/IP), the Network Virtual Terminal (NVT) Protocol and the Host Front End (HFE) Protocol. One of the major advantages of this LAN is the ability for any host or terminal to easily access other subscribers on this net or on other nets across the DDN through the use of these protocols.

SDC is planning a series of security enhancements to the MIL/INT product line that will initially add label checking and later allow full multilevel secure operation over the LAN. The initial label checking capability was to have been performed in firmware adjacent to the RF modem so that the check could be performed prior to the message being processed by software in the BIU (as shown in Fig. III-7). As development proceeded, problems have developed with this approach and it now appears that it will be necessary to perform the label checking in the main BIU program, which introduces the complexity of having trusted software to isolate this check from the rest of the untrusted BIU functions. SDC is presently using the 8086 processor for the BIU and since this is a single-state processor, it will be necessary to implement all software as trusted. Future versions of the MIL/INT product line will be implemented on the Intel 80286 processor, which supports a process structure similar to Multics. This version will readily support trusted software for both label checking and multilevel secure operation.

SDC is actively promoting the MIL/INT LAN for use in WIS and other DoD environments. The WIS Integration Contractor has contracted for an initial test bed installation of the MIL/INT LAN for possible use on WIS. A version of MIL/INT is already

being installed in a DoDIIS environment. The combination of support for the complete DoD protocol suite and the inclusion of mechanisms for label checking make this product a valuable addition to the available local area network capabilities.

7. Message Authentication Codes

Other groups are also working toward development of trusted LAN components but it appears that it will be several years before the necessary techniques are sufficiently mature to allow easy use on WIS. If, however, the LAN itself is installed in a system-high physically protected environment (e.g., protected wire line) and all resources are subjected to certain physical security constraints (to be covered in detail in the next section), then it may be possible to operate a LAN with both trusted and untrusted hosts without having to resort to E3 or trusted LAN components. An essential ingredient of this approach is the establishment of the fact that the LAN cannot modify the contents of messages sent over it by either trusted or untrusted hosts.

This same problem is faced in many situations where the integrity of information passing through untrusted components is desired. The term "spray paint," which emerged from the 1982 Air Force Summer Study on Multilevel Secure Data Base Management Systems, describes various techniques for applying an unforgetable tag to a block of information. The originator calculates the tag based on the contents of the information and appends the tag to the information before sending it through an untrusted communications medium. The recipient then repeats the tag calculation and compares it with the originator's tag. If the two are equivalent, then the recipient can have a high degree of confidence that the information was not modified during transmission. The tag calculation is usually based on a cryptographic function with a secret key known only to the originator and recipient. Typically, the information is passed through the encryption algorithm and after the block of

information has been processed, the residual value is used as the tag. Note: this technique is an integrity check, it does not provide any protection for the information being read while in transit.

This technique is being used on a number of systems to provide a means of ensuring the integrity of sensitive information. The SACDIN program will use a technique employing the National Bureau and Standards Data Encryption Standard (DES) algorithm to calculate integrity checks on its messages before they are sent to the Internet Private Line Interfaces on the Defense Data Network. The Intelligence Community is using a similar technique to ensure the integrity of data stored in a large mainframe computer in the RECON system.

The American National Standards Committee on Financial Services, X9, has published a Financial Institution Message Authentication Standard, X9.9, dated April 13, 1982, which defines a process for the computation, transmission and validation of a Message Authentication Code (MAC) using DES. The standard describes the message authentication process and the issues related to key management. This standard is being widely adopted in the financial community and implementations of it as functions on the input side of LAN interfaces may provide reasonable means of ensuring the integrity of messages passing through a LAN. The next section of this paper will discuss the application of this concept in more detail.

D. NETWORK SECURITY POLICY

The previous section described the physical, communications and computer security issues that confront a WIS LAN environment. This section will explore the security issues facing the combination of LAN and host computers. Given a reliable means of transporting messages among hosts via a LAN, what is the security policy that the LAN should enforce?

As WIS moves into the world of trusted computers, what parts of the LAN constitute the "Trusted Network Base" (the equivalent to the Trusted Computer Base defined in the Trusted Computer System Evaluation Criteria or "Orange Book")? Can WIS do anything with trusted computer systems even if it does not have a trusted LAN or a LAN with E3?

There has been much speculation on these topics and there have been several attempts to rewrite the "Orange Book" in the context of networks. But before plunging headlong into such activities, it is instructive to examine the role of various simple forms of networks in determining security policy and then projecting these results on the WIS LAN.

Let us start with the simplest case by assuming that the WIS environment will consist of a LAN similar to those described in the last section which connects a series of untrusted hosts (hosts evaluated at the D level of the evaluation criteria or hosts that have not been evaluated at all). In this situation, all hosts and the network must operate at System High and the network has no particular security relevance. It enforces no security policy since any host can communicate with any other without violating any security rules. Since there is no policy to enforce, the LAN need no security mechanisms either. This situation is very similar to the way the WIS operates today, with all the limitations on operations and performance observed earlier.

One of the goals of the WIS Modernization Program, though, is to achieve multilevel secure operation of the WIS system to overcome these limitations. As trusted computer systems become available, they will be integrated into the WIS as modules on the LAN and begin to provide substantially enhanced security capabilities. This section will explore the role the LAN must play in enforcing security separation in a WIS environment involving both trusted and untrusted computer systems.

To begin this examination, it is instructive to determine just how little trust one need place on a network. If one were to install a set of trusted hosts linked by the simplest of networks, a set of individual wires, what would the network security policy consist of and where would the security enforcement mechanisms reside? To answer these questions we need to establish a model of process-to-process communications across a network.

1. Trusted Communications over a Network with no Trusted Components

In order best to understand the relationship between the protection mechanisms provided by trusted hosts and those required in a trusted network, it is useful to explore just how much can be achieved with trusted hosts communicating over a simple network which has no trusted components. This section of the paper will examine the implications of a set of trusted and untrusted computers linked by individual wires.

It is assumed that all the resources of the hosts and communications lines are protected to a system high level (i.e., the hosts are physically protected and the communications lines are encrypted). It is also assumed that in the case of trusted hosts, the network support software of the host operating system is part of the Trusted Computing Base of that system.

2. Trusted Operating System Security Policy Model

Figure III-8 depicts the basic structure of process-to-process communications within a single host system.

All communication between processes is controlled by the Operating System (OS). In a trusted computer system, each process has associated with it a security level (SL)) which corresponds to the present security clearance level of the user on whose behalf of whom the process is operating (i.e., SL(A) = the present security clearance level of User A). In order for two processes to communicate, the following conditions, enforced by the Trusted Computing Base (TCB) of the operating system, must be true:

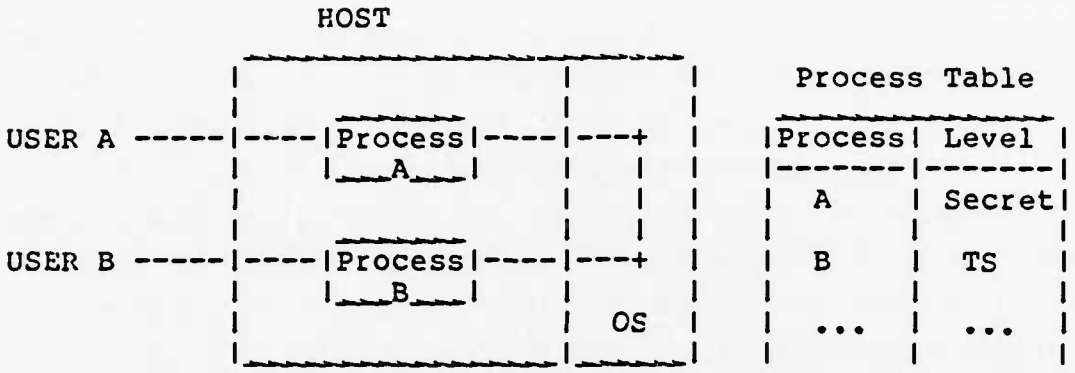

FIGURE III-8. Process-to-Process Communications within a Trusted Computer System

Process A can read information from Process B only if:

SL(A) >= SL(B) (simple security rule);

Process A can write information to Process B only if:

SL(A) <= SL(B) (* property).

In a trusted computer system of at least B2 or higher on the Trusted Computer System Evaluation Criteria, mechanisms exist for labeling all active processes and all objects as to their security level and for the establishment of a trusted path between the user and the TCB. The TCB controls access of all subjects to all objects, ensuring that the above rules are satisfied. Within the TCB of the operating system there is a Process Table of all active processes and the security levels at which they are allowed to operate. When one process attempts to establish a connection with another process, the operating system must check to make sure that the above security policy rules are followed.

3. Role of Reliable Communications

Normally process-to-process communication should take place using a two-way handshake protocol in which the receiver acknowledges receipt of the information. In the case of a trusted host where the security levels of the two processes

are not the same, such an acknowledgement process could constitute an illegal path for potentially transferring sensitive information from the higher level to the lower level. If process A is operating at the Secret level and Process B is operating at Top Secret, Process A can send information to Process B, but by the security rules stated above, Process B is not allowed to respond directly since that could violate the * property.

Within a single host this is not a particularly difficult problem. Process A can write information to Process B without an explicit acknowledgement because the process-to-process communications mechanism is assumed to be highly reliable. Once Process A has initiated the transfer, it can proceed confidently assuming that the transfer has taken place even without an acknowledgement from Process B. This highly useful simplification will not be possible when the two processes are separated by a network where the communications path is inherently unreliable.

It should be noted that this situation does not preclude passing information at one level on a trusted host to a process at a higher level on another trusted host. What must happen is that the process on the first host must establish a connection with a process operating at the same security level on the second host to pass the information with full reliable network flow control. This lower-level process on the second host can then pass the information to the higher-level process on the same host without requiring an acknowledgement. This situation does not constitute a major impediment but rather a security-related limitation which must be recognized to allow correct security operation.

4. Network Security Models

The simplest network security model (called Level 1) consists of untrusted hosts operating at the same dedicated or system-high security level and connected via an untrusted network or individual protected wires. The present WIS system is

a Level-1 network. The inadequacies of Level-1 systems have been repeatedly demonstrated in their inability to handle multiple levels of classified information. Level-1 models are included here primarily to indicate the minimum practical base upon which such systems can be built.

To overcome the deficiencies of Level-1 models, trusted computer systems like those discussed in Fig. III-8 are introduced to the simple network described above. When trusted computer systems are added, these systems become Level-2 Models as shown in Fig. III-9, which depicts two trusted hosts connected by a network consisting of individual wires.

The essential feature of a Level-2 Model is that the normal access control mechanisms in place in a B2 or higher trusted system are extended beyond protecting local process-to-process communications to handle process-to-process links across the network.

Within each computer there are processes as described previously but now, in addition, the operating system must be aware of other host computers on the network and the security levels at which they operate. The previous active process table has now been augmented with an active host table which includes information on other hosts and their security levels. There must exist a trusted means of updating this host-security level information but since on a host-by-host basis this information is very stable, this update could be performed in the simplest case by a periodic manual Security Officer table update operation.

There are a number of examples which illustrate the basic access control checking flow of this model. Consider when Process A, operating at the Secret level, attempts to communicate, System 1 receives the request and checks to be sure that the level of A is within the range of Host 2. If it is not, the request is denied. If it is, OS1 passes the identity and security level of A to OS2 which checks if the level of A

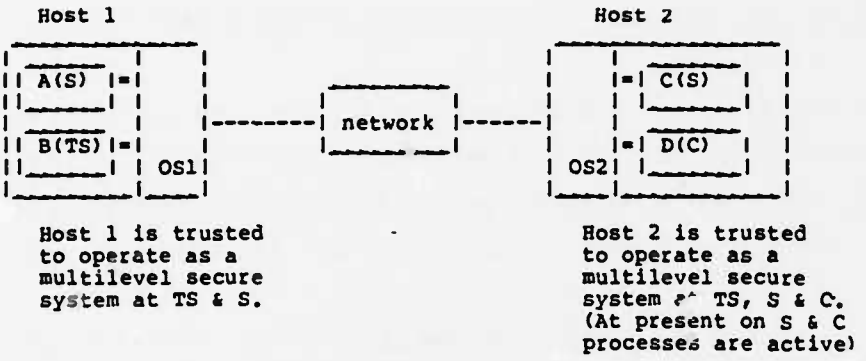

Operating System Process Tables:

Host 1

Processes	Hosts
A at level S	2 at levels TS-C
B at level TS	...
...	

Host 2

Processes	Hosts
C at level S	1 at TS & S
D at level C	...
...	

Note 1: OS1 and OS2 must have individual authentication codes that they use to ensure against spoofing from other hosts on the network.

Note 2: Processes on different hosts (or the operating systems that support them) may employ cryptographic checksums on the data messages before sending them to the host. These checksums ensure against data modification while in transit.

FIGURE III-9. Level-2 Network Security Model (Both Hosts Trusted)

AD-A152 996 PROTOCOLS AND SECURITY IN THE WWMCCS (WORLDWIDE MILITARY COMMAND AND CONT..(U) INSTITUTE FOR DEFENSE ANALYSES ALEXANDRIA VA T C BARTEE ET AL. NOV 84 2/2

UNCLASSIFIED IDA-P-1796 IDA/HQ-84-29059 MDA903-84-C-0031 F/G 17/2 NL

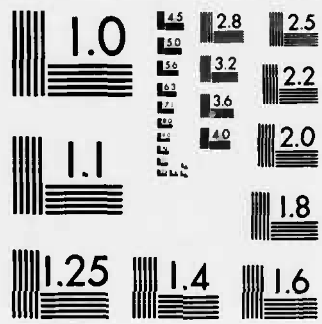

equals the level of C. If it is, the connection is extended to the two processes. If not, OS2 informs OS1 that the connection is invalid.

A second case to consider is if Process B, operating at Top Secret, attempts to connect with Process C, operating at Secret. OS1 checks that the level of Process B is within the range of Host 2. Since it is not, OS1 immediately rejects the request.

If C wanted to send information to B (where SL(B) = TS and SL(C) = S), C would have to establish a connection with a Secret process on Host 1 to pass the information, using a fully reliable protocol with acknowledgements, and then this process would pass the information to B, as provided by the Bell-LaPadula model on the same host.

The third case is when Process A, operating at Secret, attempts to connect with Process D, operating at Confidential. OS1 checks if the level of Process A is within the range of Host 2. It is, so OS1 passes the identify of Process A and its security level to OS2 which checks if the level of A is equal to the level of Process D. It is not, so OS2 informs OSL that the connection is invalid.

The Level-2 model represents the simplest structure involving trusted computers. It is patterned after the present structure in use by hosts on many networks. Much of the simplicity is achieved by not requiring (or allowing) host operating systems to know anything about processes that may exist on other computers on the network. A more complex model could be constructed wherein the status of foreign processes (existing on remote hosts) could be included in a local host's processor table. Such a structure has the advantage of allowing a local host to make all access decisions based on local information, eliminating the steps to query the remote host. But such a structure poses difficult problems of maintaining data bases of remote host processes in a trusted manner. In contrast to the

stability of the security levels associated with a computer system, the status and or security level of individual processes must be expected to change very dynamically. The problems inherent in maintaining such a complex data base are not worth the small enhancement in performance that might be achieved.

Figure III-9 assumed that both operating systems were trusted so that they could rely on each others' access control mechanisms. The WIS LAN environment will contain systems which are either totally untrusted or cannot be trusted to the same level. Figure III-10 is an augmented version of Fig. III-9, with an untrusted host operating at the Network System-High Top Secret level attached to the net (still assumed to be just individual wires connecting the hosts). The same caveats associated with the network in Fig. III-9 apply here.

Note the addition of Host 3 at a single security level to the host tables in Hosts 1 and 2. The Process Table in Host 3 knows about Hosts 1 and 2 since it is capable of communicating with them at the Top Secret level, but this table contains no security label information since the host is not trusted. As shown below, Hosts 1 and 2 will accept connections from Host 3, recognizing that it is a system-high Top Secret system and that all process-to-process communications with it must be held at the Top Secret level.

There are a number of additional cases that must be examined for the configuration shown in Fig. III-10. The first is when Process E, operating at Top Secret, requests a connection with Process B, also operating at Top Secret. OS3, which is not trusted and does not have a security level associated with its Process-Host table, establishes a connection with OS1 which checks if the level of Process B is equal to the system-high level of Host 3. If it is, the connection is passed to Process B. If not, OS1 rejects the connection.

In the next case, Process A, operating at the Secret level, requests a connection with Process F on Host 3. OS1 checks if

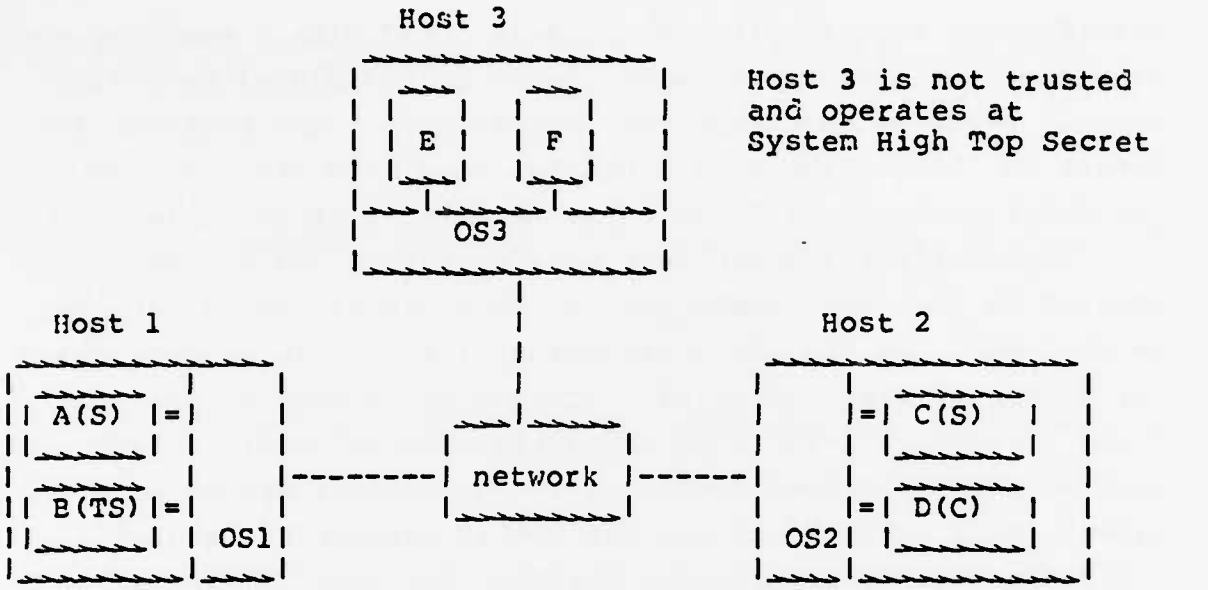

Operating System Process Tables:

Host 1

Processes	Hosts
A at level S	2 at levels TS&C
B at level TS	3 at level TS
...	...

Host 2

Processes	Hosts
C at level S	1 at TS & S
D at level C	3 at TS
...	...

Host 3 (since this host is not trusted, its process table does not contain any security level information).

Processes	Hosts
E	1
F	2
...	...

FIGURE III-10. Level-2 Network Security Model with an Untrusted Host

the level of Process A is equal to the system-high level of Host 3. It is not, so OS1 rejects the connection. The third case involves Process E on Host 3 attempting a connection with Process D on Host 2. OS3 initiates a connection with OS2. OS2 checks if the level of Process D is equal to the system-high level of Host 3. It is not so OS2 rejects the connection.

At this point we have described a network of trusted and untrusted computers that are able to communicate in useful, practical ways using a network with no trusted components, and enforcing no security policy or access control mechanisms. The next step is to understand the limitations of such a network. Figures III-9 and III-10 consist of host computers that are physically protected to a network system-high level (Top Secret is assumed in both cases). Host 2 only has processes currently running at the Secret and Confidential level but it must be operated in a Top Secret system-high environment. Host 3 in Fig. III-10 is an untrusted host which must also operate at Top Secret System High.

The relationship between the security levels at which hosts can operate on this simple network is shown in Fig. III-11. The network itself and all computers attached must be physically protected to the same system high level. Untrusted hosts can only operate at that level. Trusted hosts may be allowed to operate over a range with a maximum at the network system high. The range over which they can operate is dictated by the degree of trust in the system. In Fig. III-11, Host 1 is trusted to the "A1" level and operates over the range Top Secret to Confidential (range refers to levels of clearances of users).

Host 2 is trusted to the "B2" level and is allowed to operate with users cleared to either Top Secret or Secret. Host 3 is untrusted and operates only at Top Secret. Host 4 operates over the range Secret and Confidential on a Top Secret System-High network. Since it is possible that Host 4 will

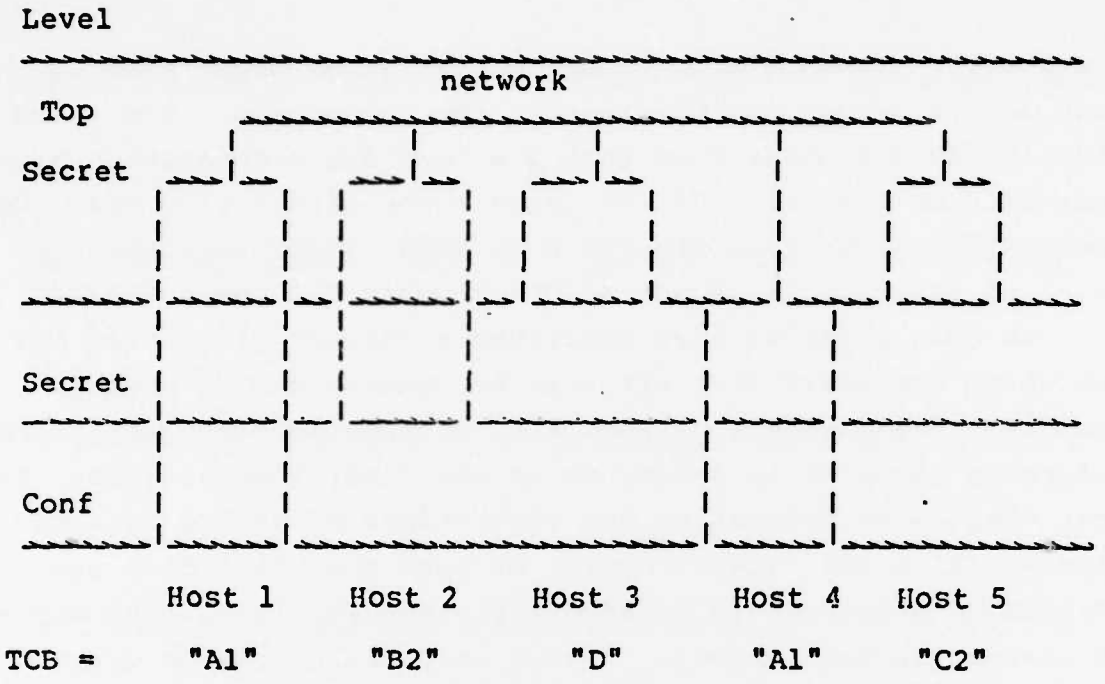

FIGURE III-11. Trusted and Untrusted Computers on a System-High Network

receive a Top Secret message from an untrusted host, the system must be capable of operating over the full range of Top Secret to Confidential, therefore Host 4 must be trusted to the "A1" level. Host 5 is trusted to the "C2" level which provides only discretionary access controls. Since no provision for mandatory labeling is contained with the "C2" level, Host 5 must operate at the Top Secret level.

Trusted hosts must have the highest point of their operating ranges at the network system high since they can receive system-high messages from untrusted hosts. Similarly, since untrusted hosts are not capable of protecting the security labels of information, they must be operated only at system high.

Note: the ranges shown in Fig. III-11 over which hosts operate are arbitrary selections. In an actual system the

Designated Approving Authority or System Security Officer would determine the degree of trust needed for a particular situation.

5. Trusted Computers on More Sophisticated Networks

As depicted in Fig. III-11, it is possible to build a network with a full range of trusted and untrusted computers in which the network itself (still just a set of wires) has no trusted components. As long as the entire system is protected physically from external attack, the network need not be trusted since there are no network components that could alter the trusted labels produced by the trusted hosts or otherwise disrupt service.

But networks composed of sets of wires between hosts are not very practical. The local environment of the WIS will employ LANs and the total WIS environment will require the LANs to be linked to each other by the Defense Data Network (WIN). This is a much more sophisticated environment than that described above and we must determine the feasibility of extending the above ideas to more complex networks.

As described in the previous section, trusted Bus Interface Units provide an excellent means of ensuring the integrity of messages being sent over a LAN. However, trusted BIUs are still not readily available. The use of Message Authentication Checks (MACs) over a local area network is a viable means of ensuring the integrity of data being sent over a physically protected but untrusted medium. If we were to apply MACs at the host interface to the LAN, we could have considerable confidence that the LAN is not able to modify the contents of messages without detection.

Each WIS LAN will have a gateway connection to other LANs via the DDN. As presently configured the WIN is a system-high net with only WWMCCS hosts attached. The present WWMCCS hosts are all untrusted and therefore run at System High. Once the WIS LANs are capable of handling both trusted and untrusted

hosts as described above, provisions will have to be made to protect the integrity of messages passing over the LANs and or DDN.

Figure III-11 depicted our simple network composed of individual wires linking both trusted and untrusted hosts. When this simple network was replaced by a LAN it was necessary to add either trusted BIU devices or host interface MACs to ensure the integrity of messages passing over the network. When these LANs are linked by a network such as WIN it is necessary to introduce either trusted gateways and trusted packet switches (which are not practical in the near future) or extend the coverage of the MACs between hosts on multiple LANs. If more than just WWMCCS sites are attached to the WIN, end-to-end encryption devices such as the IPLI or Blacker devices described in the preceding section will need to be added. This situation is depicted in Fig. III-12.

A MAC must be calculated for each message entering the LAN, either in trusted software in the host device or in a special interface box on the host or LAN. The interface with the WIN gateway must be a special back-to-back MAC (probably using different encryption keys) to authenticate messages leaving the LAN to apply MAC protection to the message as it enters the gateway and WIN. At the receiving WIN connection a similar back-to-back MAC authenticates the incoming WIN message and applies a new local MAC for entering the message into Site B's LAN. (It may be possible to bypass the MAC device on the gateway connection and have the LAN MAC apply to the message as it flows across the WIN and into Site B's LAN. This approach would require that all LANs which wish to communicate must use the same encryption key in calculating the MAC, a situation which is not attractive either from a security or operational point of view.)

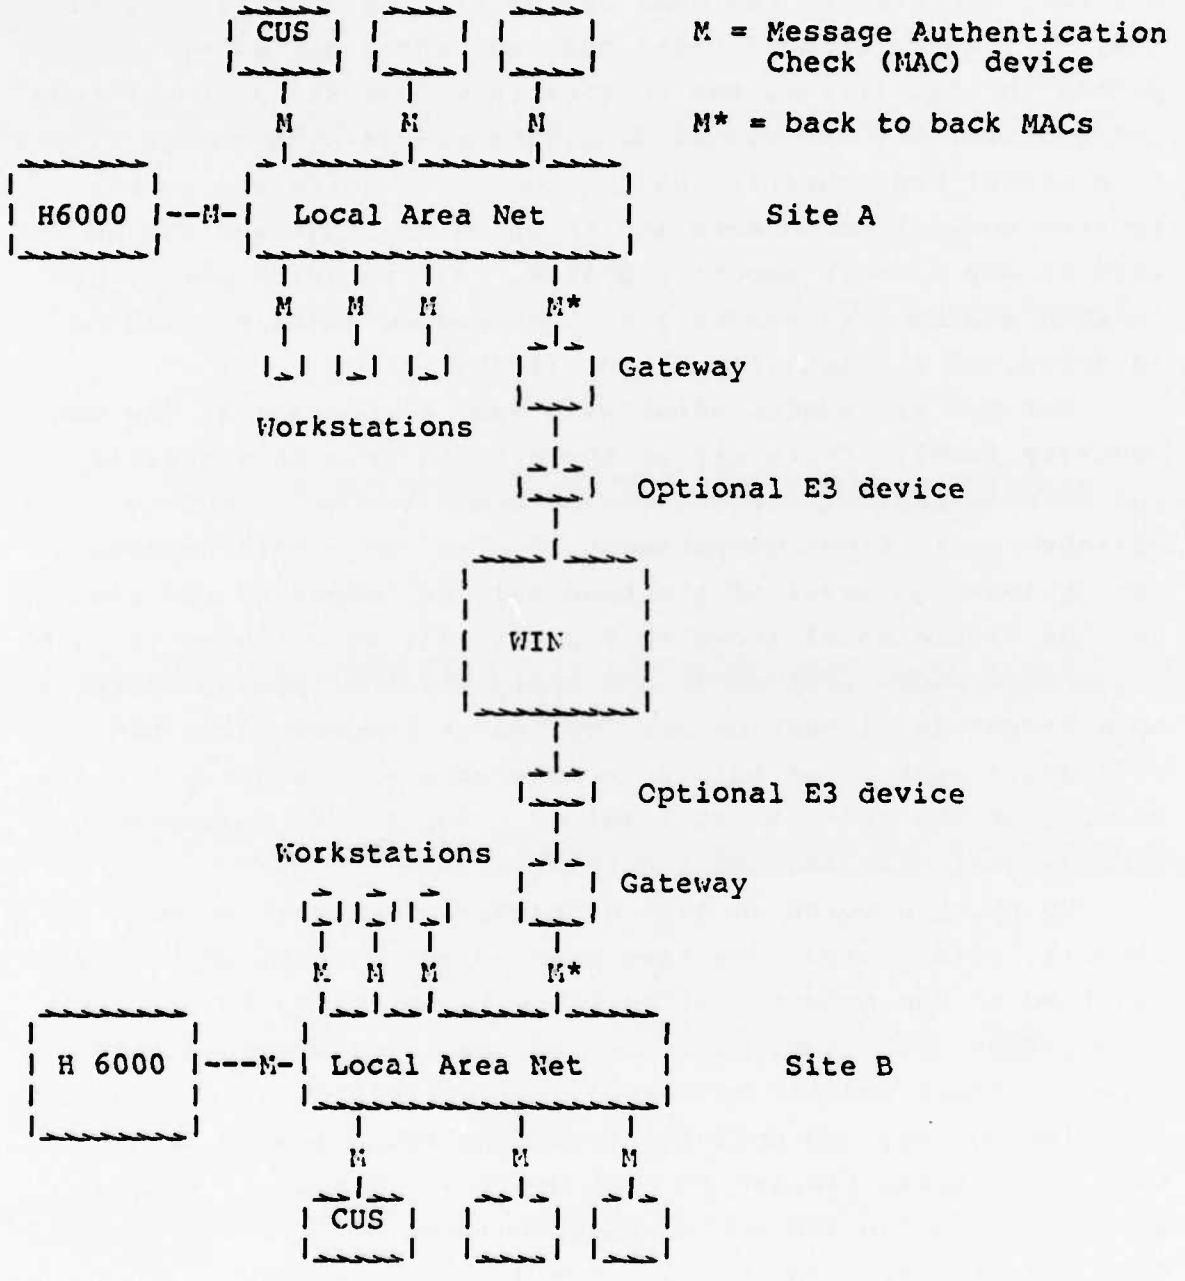

FIGURE III-12. WIS LANs and WIN with Message Authentication Checks

6. WIS LANs Operating at Different Security Levels

The above discussion described the situation where all WIS LANs operated at the same system-high security level, as they do today. Trusted hosts had been added but as was depicted in Fig. III-11, the network (now consisting of the LANs and the WIN, covered by the MACs) are all at System High. This is a useful configuration which does not require any special trusted network components and in which the LANs and WIN do not enforce any network security policy. All security policy decisions are handled entirely in the trusted hosts themselves as described in Figs. III-9 and III-10.

But not all WWMCCS sites will want to operate at the same security level. Certainly as the systems grow in complexity and take on new functions, some of the sites will want to operate at different compartment levels. Once this happens the system-high level of the LANs will no longer be the same and the simple model shown in Fig. III-11, as extended through the use of MACs, will no longer apply. Yet processes operating on a Secret-level host on one Top Secret Compartment A LAN will still want to be able to communicate with other processes running at the same Secret level on a Top Secret Compartment B LAN. How will we accommodate this?

Up to this point we have no required explicit network security policy and there have been no enforcement mechanisms required in the network (since there is no policy to enforce). This latest situation, however, requires the introduction of a network access control mechanism, the definition of a network security policy, and policy enforcement mechanisms. We now have a situation similar to that depicted in Fig. III-11 but with the networks and system-high hosts at different levels. This case is depicted in Fig. III-13.

For the sake of simplicity, the two LANs are shown with only two hosts each, one running at both TS Compartmented and Secret and the other only at the Secret level. (It should be

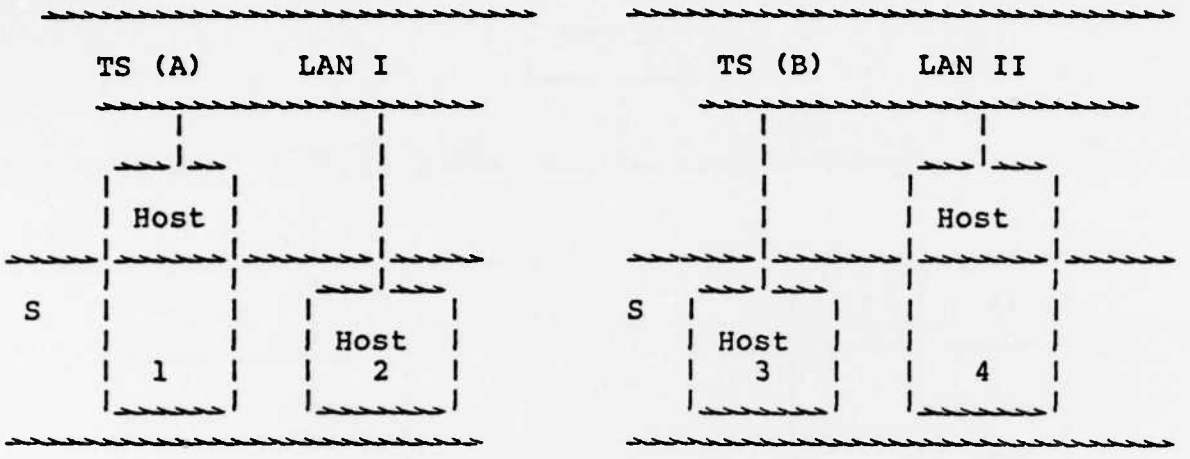

FIGURE III-13. Two WIS LANs at Different System-High Levels

recalled that the host running at the Secret level must be capable of running at both Top Secret Compartmented and Secret in order for the conditions described in Fig. III-11 to apply.) It is not possible for processes at the TS Compartmented levels to communicate since that would violate the separation of compartments. It could be possible (and will be expected) that processes running only at Secret on Hosts 2 and 3 should be able to communicate.

To achieve this level of operation it will be necessary to resort to some form of network access control to determine whether processes on specific hosts can communicate. The IPLI with its static host connection tables is not sufficient for this task but the dynamic checking envisioned in the Blacker system may provide the needed control. Figure III-14 depicts the two LANs in Fig. III-13 connected over the WIN with the Blacker E3 devices as described in the previous section. The access controller/key distribution center function provides the crucial network access control mechanism to enforce connections across LANs operating at different Network System-High Levels.

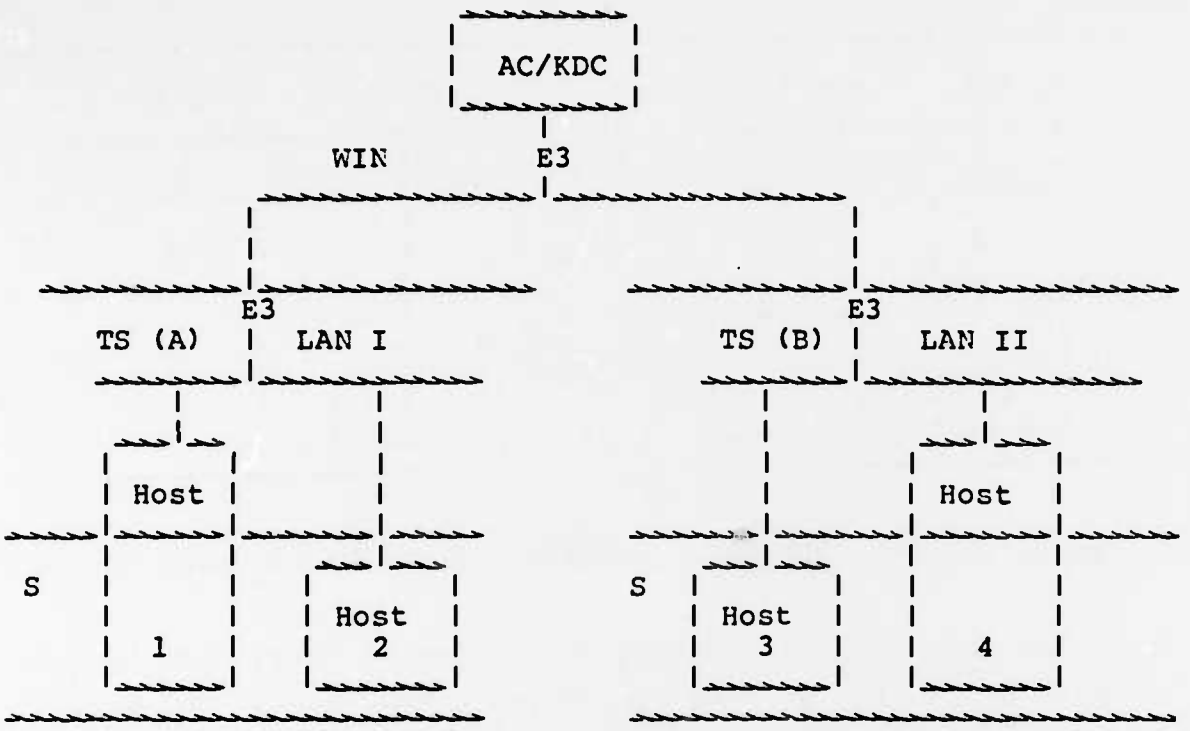

FIGURE III-14. Two WIS LANs at Different System-High Levels Connected with Blacker End-to-End Encryption Devices

Assuming that Hosts 2 and 3 are operating only at the Secret level and that this fact is known to the Blacker access controller, an attempt by a process on Host 2 to connect to a process on Host 3 will be allowed since both are limited to operating at the same level. If the Blacker access controller could communicate with Host 1 in a trusted manner, then this same procedure would allow a process running at the Secret level in Host 1 to communicate with a similar Secret process on Host 3. However, processes running at the network system high would not be allowed to communicate since the network system-high levels are different.

7. WIS Hosts at Different Levels on the Same LAN

The first example we explored was a single LAN with trusted and untrusted hosts all operating at the same Network System-High level. Trusted hosts were allowed to operate over a range that extended from that Network System High down to some lower level of sensitivity, depending upon their level of trust (e.g., hosts trusted to "A1" might range from Top Secret to Confidential while hosts trusted to "B2" might range from Top Secret to Secret). This network model had the significant advantage that there was no explicit network security policy and therefore no Network Reference Monitor or Trusted Network Base. All that was required was a trusted path across the physically protected LAN, a path provided by the use of Message Authentication Checks, for example. All security policy enforcement was performed by the hosts themselves.

The second example we explored was where two or more such LAN environments operating at different Network System-High levels were connected via a wide area network such as the WIN. In this case, a network security policy was required to check if the levels of the processes wishing to communicate were equivalent. Mediation at the network level was required since the two Network System-High levels were not comparable. A network policy enforcement mechanism was required to mediate requests across these LANs. There are no trusted mechanisms on the WIN that can perform this mediation. Sometime in the future though, it is likely that a form of end-to-end encryption (E3) similar to that proposed by the Blacker Program will be included in the WIN. In the case explored above, the E3 Access Control mechanism performed the mediation between processes on LANs operating at different Network System-High levels. The E3 system then becomes the Network Reference Monitor or Trusted Network Base.

However, there are other situations that will arise in the WIS for which neither of these network models are sufficient.

In both these cases, all untrusted hosts were required to operate at the Network System High level. This includes the Honeywell 6000 computers and any other functional modules that are not implemented on trusted computers. In a typical environment where the H6000s operate at Top Secret, this restriction would preclude any untrusted hosts at the Secret level.

To provide for this situation it will be necessary to resort to some form of trust within the LAN itself. It will be necessary for the LAN to mediate access between the security levels of the various devices on the network since those devices cannot reliably determine the level at which other devices operate. This mediation can be performed by E3 devices on the LAN or by Trusted BIU devices, as described in the previous section. Under these circumstances, trusted and untrusted hosts operating at any security level can be attached to the LAN. The security level at which an untrusted host operates (or security levels in the case of a trusted host) are known to the LAN Trusted Network Base and attempts to communicate between processes on different hosts must be mediated by the TNB.

If E3 is used on the LAN, then the TNB consists of the E3 devices at each LAN port and the Access Controller and Key Distribution Center that comprise the entire E3 process. The Access Controller has prior knowledge of the security level at which processes on specific hosts are allowed to operate. When a process on a specific host requests a connection with a process on another host, the Access Controller checks to be sure the levels of the two processes are equivalent (reliable communications across a network still requires that the processes be equivalent). If they are, the AC allows a key to be distributed to each E3 device so that the connection can be established. Depending upon the sophistication of the AC, this approach can enforce discretionary (need-to-know) access controls as well as mandatory controls. However, discretionary controls using

access control lists, for example, will be difficult to maintain on a process-by-process basis across the network. A distinct advantage of this approach is that no portion of the LAN need be trusted.

If Trusted BIUs are employed, then the trusted portions of the BIUs define the TNB. In this case typically, the trusted BIU is aware of the level(s) at which the host attached to it is allowed to operate. When a process on one host attempts to communicate with a process on another host, the originating BIU passes the security level of the requesting process to the destination BIU which checks to make sure that the security levels are equivalent. If they are, the two BIUs allow the connection to take place. This approach works well for mandatory access controls where every process has an explicit label associated with it. As in the above E3 case, discretionary controls are harder to apply since they require maintenance of extensive access control lists in the BIUs.

This third network model where untrusted hosts operating at different security levels are attached to the same LAN represents the most common and desired environment for the modernized WIS. It provides the richest set of possible operating scenarios with the minimum operational limitations. However, as was pointed out in the previous section, neither End-to-End Encryption of the type needed for this environment nor Trusted BIUs are available today and it is likely to be at least four to six years before such facilities are available.

Without some form of Trusted Network Base to enforce the network security policy described above, it will be necessary to operate WIS LANs according to the restrictions contained in the first two examples above where all untrusted hosts operate at the Network System-High level and all trusted hosts operate at a range extending downward from that Network System High, depending upon their degree of trust.

E. ACCESS CONTROL PROCEDURES FOR WIS

The previous section explored the requirements for network security policy in a multiple WIS LAN environment connected by the WIN. Situations requiring an explicit Trusted Network Base were examined and instances where individual LANs could operate without any network security policy or TNB were described. Since it will be some time before either End-to-End Encryption or Trusted Bus Interface Units are readily available, these alternative forms of LAN operation may define the mode of operation of at least early versions of the WIS LANs.

The previous section made a number of simplifying assumptions so as to not overcomplicate the analysis being performed. It was assumed that all entities that connected to the LAN were host computers and that any terminals to which users were connected were attached directly to these hosts. This is not a particularly limiting assumption since most new WIS "terminals" will be at least of the personal computer form with considerable processing power available locally. Such devices should be considered hosts rather than simple terminals and should have fully functional host interfaces to the LAN. The other major assumption was that user authentication was performed by these hosts and the LAN was not involved. The models discussed concerned existing processes that wished to communicate across the network. The issue of how these processes came into existence was a responsibility of the hosts themselves.

The set of issues to be explored in this section includes the development of means of evolving the user authentication mechanisms in use today on the Honeywell 6000 computers to the modernized WIS environment. The form of this evolution will be impacted by the type of trusted or untrusted LAN that is employed in the WIS but it is also affected by many other issues which will be examined here.

The present structure of the WIS is shown in Fig. III-15. All user log-in and authentication takes place within the H6000s themselves. Valid user names and passwords are required to determine the user's access to sensitive information. When users wish to perform remote log-ins across the network, the authentication processes on the local host forwards the user authentication information to the remote host as part of the network connection protocol interactions. All log-in and accounting functions are performed in the H6000s themselves and since the H6000s are all compatible (at the operating system level at least), they trust each other to handle user authentication information properly.

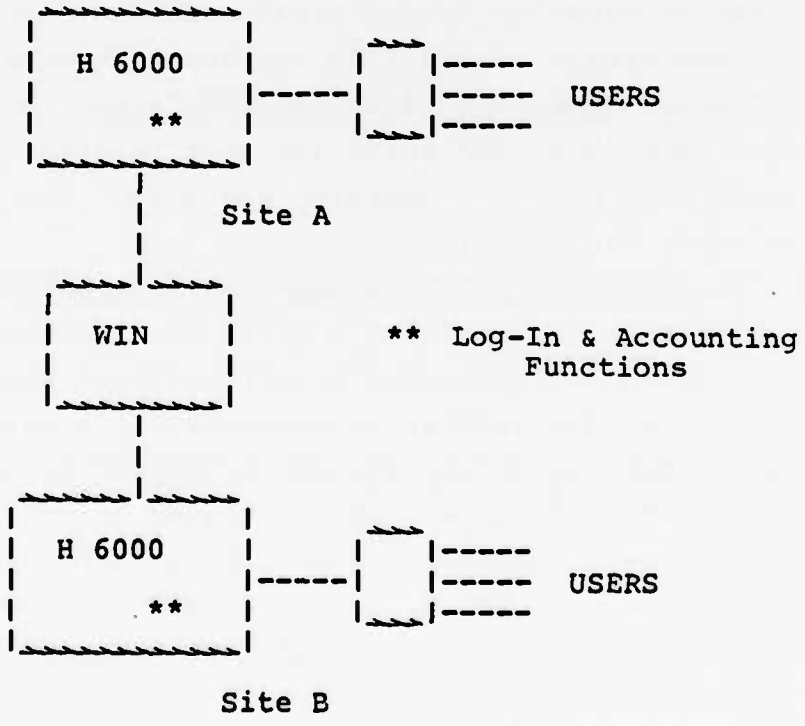

FIGURE III-15. Present WWMCCS Environment

When the new WIS LAN structure is in place, the issue of user log-in and authentication becomes much more complex. The new structure is represented in Fig. III-16, which shows the H6000s, LANs, special functional modules such as the Common User Support (CUS) system and an array of work stations at two WIS sites connected via the WIN. In general, there are two types of user authentication, each with several variations: centralized and distributed. Centralized authentication involves special hosts on the network, which contain an authorized user data base that is checked before a user is allowed access to any resource on the LAN. Distributed authentication can take place at the point of origin of the user access (e.g., the work station) or at the point where the processing is to take place (e.g., the H6000 or other functional module).

Access Control Procedures for the present WWMCCS system are similar to those in conventional multiuser mainframe computers of the early- to mid-1970 vintage. Terminals are attached through front-end processors to the mainframe and user log-in is handled by special operating system functions that require a proper user name and corresponding password. The procedure is depicted in Fig. III-17.

The user signals the GCOS operating system to start the log-in process (1) by striking a Break or Attention key. The operating system requests that a user name and password be supplied. Once this process is completed, the user is allowed to invoke specific programs (2) on the local host and to process data to which he or she is authorized access. Among the programs that may be authorized are those necessary to establish connections over the WIN to other WWMCCS hosts. In this case the user invokes the Telnet and NCP processes (3) to establish a connection with the remote host. The remote host invokes its log-in process (4) to confirm the user's access privileges. If the user is authorized to access specific information on the remote host (5), this second log-in process

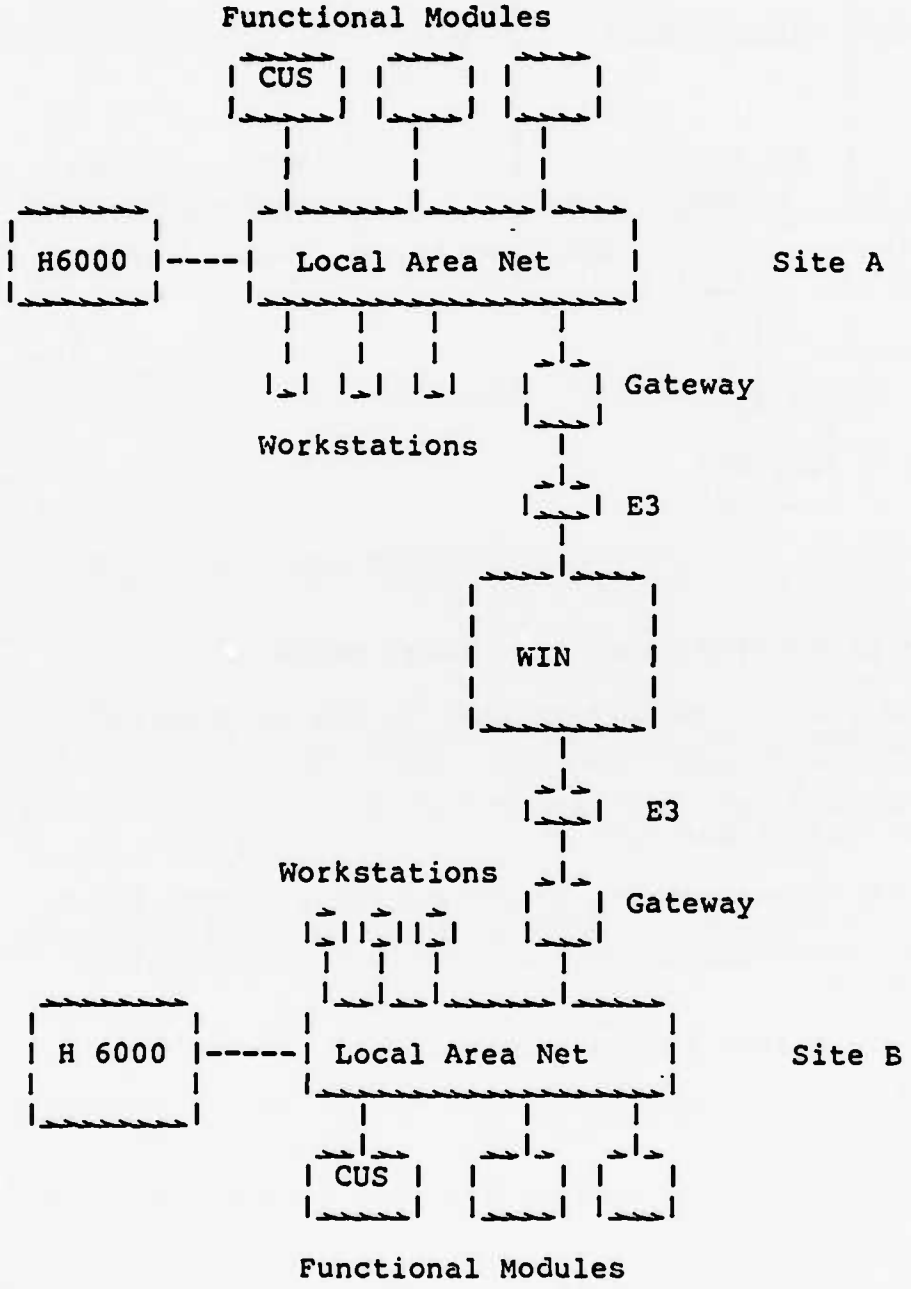

FIGURE III-16. WIS Modernization Environment--User Authentication in the Present WWMCCS ADP System.

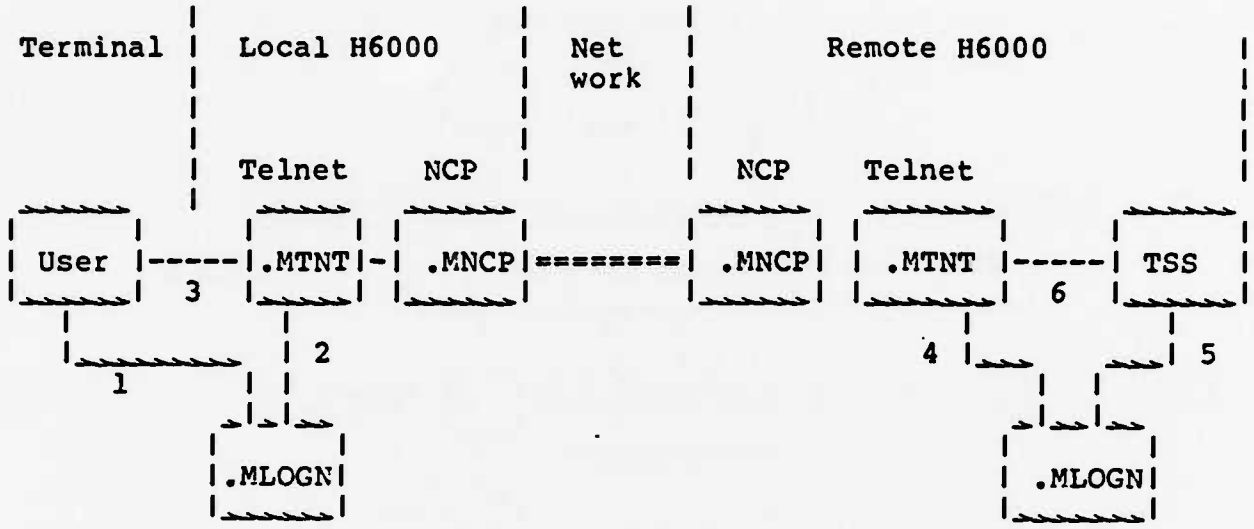

Figure courtesy of CSC

1. User at terminal logs into local H6000.
2. Login process authorizes user to initiate specific processes.
3. Authorized user invokes network connection processes (User Telnet and NCP).
4. Remote Server network processes require user login.
5. Authorized remote user allowed to invoke specific processes.
6. User initiates remote processing.

FIGURE III-17. Network Security Flow--Terminal User

controls the nature and extent of that authorization. Finally, the user initiates the desired remote process (6). In this structure, each host is in control of its own authentication processes and it is necessary for remote users to be added formally to the authorized user tables before they log-in.

1. Untrusted Host Environment

The following section describes possible methods of user authentication on a WIS LAN where there are only untrusted host computers such as the H6000s. The LAN, hosts and the WIN connections are all assumed to operate at a Network Top Secret system-high manner (exactly as they do in WWMCCS today). In this environment, user authentication is needed to provide a means for auditing user actions but will not be trusted to grant access to multiple levels of sensitive data. User authentication in the presence of trusted hosts will be discussed later.

2. Mainframe Log-In

Since this mode of LAN operation does not require any more trust than is available in the present H6000 computers, as shown in Fig. III-15, the simplest form of user authentication may be a direct extension of the present mainframe computer log-in mechanisms, exactly as is done now. In this context an H6000 computer and any of the functional module systems are considered mainframes and the work stations are considered to be terminal devices (even though useful work can be performed on a standalone work station). Users at work stations would gain access to the network without logging in but as soon as they attempted to access a particular mainframe, they would be forced to go through that mainframe's log-in process.

This approach has the strong advantage that it requires little if any change in the existing H6000 hosts, depending on how the H6000s interface to the LAN. The user, after requesting a connection to a particular H6000, would go through exactly the same user log-in required on the present system. If use of a remote host is required, the same network connection and remote would be required.

There are many problems with this approach. While the present H6000 log-in process is well-understood by the user and implemented uniformly on the H6000s, it may not be optimum in a local area network multiple computer environment.

This procedure may be difficult to implement on hosts other than H6000s, with the result that the log-in process would appear different on different hosts, a most unfortunate human interface characteristic. The interactions between different types of computers on the same LAN may be sufficiently constrained that it will be difficult to effect the remote log-in process in a reasonable way.

While it is clear that a mainframe log-in mechanism is possible, perhaps the most serious drawback to this approach is that it represents a large step backward (or at least sideways) from the other major objectives of the WIS. The notion of workstations connecting to functional modules which perform specific tasks on behalf of the user, holds as a basic premise that the user should not have to know about what process is being run on what computer on which network. The goal is to achieve an overall system in which the user requests that a service be performed and devotes full concentration to understanding the results of the computations. If the user must log-in to every host computer on the LAN that may be involved in the computation, then he or she will be distracted significantly from the task to be performed while going through various log-in steps. Even if this log-in issue is properly handled, there remains a question of whether the overall WIS system will be able to achieve the objective of isolating the user from awareness of the various computers that may be running portions of his or her task. But it is clear that this goal cannot be achieved if simple mainframe log-in is required.

Another issue of concern with this approach is whether users will be required to use different log-in names and passwords on different hosts. If this were to be required, users would have a very difficult time handling more than three or four functions on different hosts and there would be a strong tendency to write down various log-in passwords, or worse, put

them into programs on one machine to "ease" the burden of remembering how to log into other hosts. If on the other hand, a user employs the same name and password on all hosts to which he or she may connect, first, that information is now exposed to many opportunities for compromise on the various hosts (the user will tend to register on more hosts than normally needed just in case a special situation arises) and second, if a compromise occurs it will be difficult and cumbersome to correct since the user probably will not remember all the hosts on which he or she is registered, and there is no central mechanism to keep track of this information. Once a password is compromised, it must be changed on every host or that system remains vulnerable.

One unacceptable method around this issue would be to have every user registered for every host on the entire LAN. This would make it relatively easy to correct compromises of user passwords since a broadcast message could be sent to every host with the update. Unfortunately, this approach does not generalize well for remote users on other LANs, nor is it practical for large sites with hundreds of users. It also violates the security premise that users should be isolated to only those machines for which access is required.

3. Work Station Log-In

A second form of log-in in a LAN environment where only untrusted hosts are present involves requiring the user to log-in on the local work station prior to performing any tasks on the work station or accessing any remote mainframes. This approach has several strong advantages from a security perspective. It eliminates the problem in the previous example where users could do considerable work on workstations without any authentication or audit control. It also provides the opportunity at least for uniform log-in and authentication across all workstations on a LAN. If the mainframe systems could trust the workstation log-in, then a second log-in at the mainframe would

not be needed, but since this is an untrusted environment, log-in mechanisms similar to those presently in use at each mainframe will still be required.

There are major issues with this approach which decrease its attractiveness. The first is the problem of establishing the data base of authorized users that could access any particular workstation. The list at any single workstation would want to be some subset of the total list for a WIS site, but what subset? And then once the subset has been sent to the workstation, how is it protected? Having user authorization and password information spread among all the workstations throughout a site is not a viable security procedure.

4. Network Log-In

There are several examples available of network log-in procedures that are quite effective at providing uniform single log-in mechanisms when fully implemented throughout the system. One which is in use on several operational networks at the present time involves the use of special terminal access devices called PADs to interface remote terminals to a network and several redundant Name Authentication Servers (NAS) which contain the user name, password and type of service authorized for each user (See Fig. III-18).

The user requests service from the network by either striking an attention key if the terminal is on a dedicated line or by dialing up the PAD. The PAD immediately requests that the user supply a Username and password. The password is encrypted using the DES algorithm in a one-way encryption mode and the username and encrypted password are sent over the network to the NAS. The NAS is basically a large data base server system which uses the username and encrypted password as the key to search for the complete user's entry in its data base. Once the entry is found, the entire record is returned to the PAD including whatever privileges the user may have. If the user is only allowed normal log-in privileges to certain

mainframe computers, the PAD then establishes a special constrained connection to one of those computers and an automatic log-in process takes place. The user goes through a single log-in process and has ready access to whatever resources are authorized on the network.

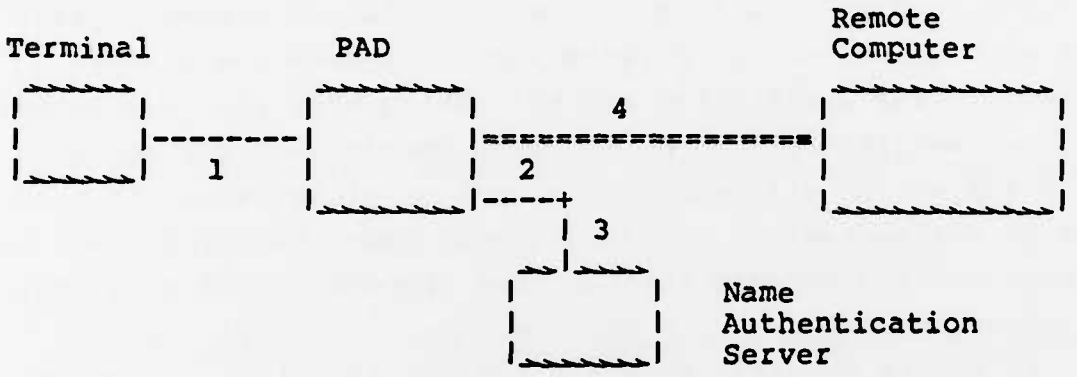

1. User requests network service, PAD demands Username and Password
2. PAD encrypts password and sends with username to NAS.
3. NAS looks up user privilege information based on user name and encrypted password and passes back to PAD.
4. PAD determines user privileges and establishes automatic connection with authorized remote computer.

FIGURE III-18. PAD-NAS User Authentication Process

This technique relies upon the correct functioning of the PAD, a device under the control of the network which has very limited functionality and a carefully controlled software environment. For this process to work on a WIS LAN the workstations would have to duplicate the PAD log-in acceptance function as an initial process which could not be bypassed.

They would then pass the user log-in information to the NAS for validation. The NAS could pass either the full user record as it now does or a subset listing authorized points of access for that user back to the workstation which would proceed to allow the user to perform authorized operations either locally at the workstation or at remote computers over the LAN.

The problem with this approach is that it places a great deal of reliance on the workstation to perform the checks properly. The workstation may be nothing more than a personal computer and it resides in the local user's physical space so it is subject to all manner of tampering either with the hardware or the software. To have to rely upon the workstation to perform all the PAD-NAS interactions properly would pose serious security risks.

It may be possible to place the PAD functions in the LAN BIU which, while also in the user's space, is hopefully much less vulnerable to tampering either from a hardware or software perspective. The user could presumably use the workstation without formally logging in for local functions but the BIU would force a proper log-in before any communications could take place over the network. Relying on the workstation to supply the user name and password correctly remains a problem with this approach.

5. TAC Access Control System (TACACS)

The Defense Data Network has installed a variation on the above user authentication process which is called TACACS. This technique is implemented in the DDN Terminal Access Controllers (TAC) as an access control means for dial-up ports on the DDN. An essential ingredient of this approach is the Self Authenticating Password (SAP), which is based on a DES encryption of the user's identification and a version/group number (Figure III-19). When a user registers for access to the network, he or she is sent, via U.S. Mail, a User ID and SAP. In the first phase of TACACS implementation (Figure III-20),

when a user attempts to access the network the TAC requests that user to supply both items. The TAC then passes them to its local IMP where the DES algorithm and key are stored in microcode. The IMP calculates the SAP from the User ID and if it matches the user supplied SAP, notifies the TAC that the user is authentic.

Self Authenticating Password

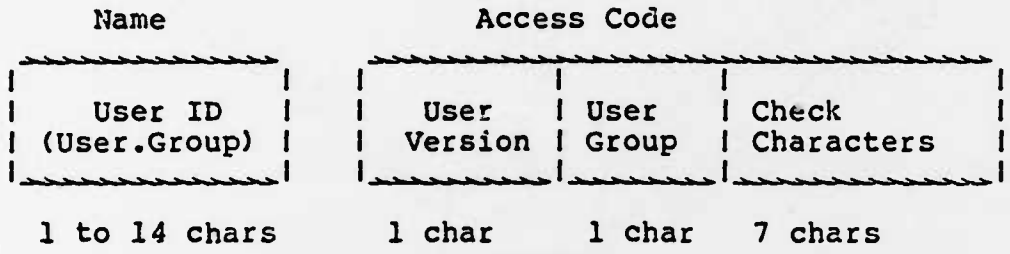

FIGURE III-19. TAC Access Control System

This system provides a good a means of user authentication which is very difficult to guess and, as long as the user takes reasonable precautions with his or her SAP, resistant to compromise. Audit information is passed periodically from the TAC to a Data Base Host. This information provides specific user identification of TAC user for the first time in the history of the ARPANET. There is always the possibility of loss of a SAP so there is provision for a "hot list" of disallowed passwords in the TAC. A user who reports a lost SAP will be issued a new one based on the same user ID but using a new user version number. Since a change of only one character in the sequence used to calculate the check characters will result in a totally new set of check characters, this

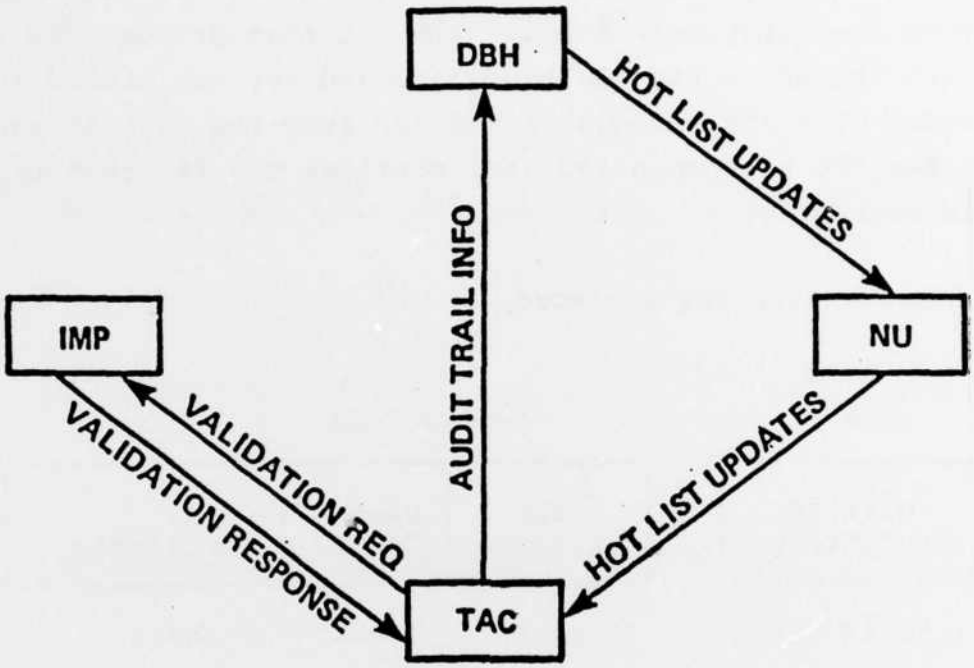

FIGURE III-20. TAC Access Control

approach combined with the "hot list" of bad SAPs provides a practical means of keeping the access lists up to date. When needed, whole groups of users can be issued new SAPs by changing the User Group Version Number.

Later phases of TACACS provide for Log-in Hosts (Figure III-21) to perform the user validation process. The same user ID and SAP will be employed but the Log-in Host can handle the "hot list" problem more efficiently. Eventually, an approach using a Personal Authentication System is planned (Figure III-22). This system involves the use of a credit-card-like "Smart-card" in conjunction with the user's terminal. Upon attempting to log-in to the network, the user will be sent a "challenge" character string by the Name Authentication Server. The user will enter this challenge, together with his or her password,

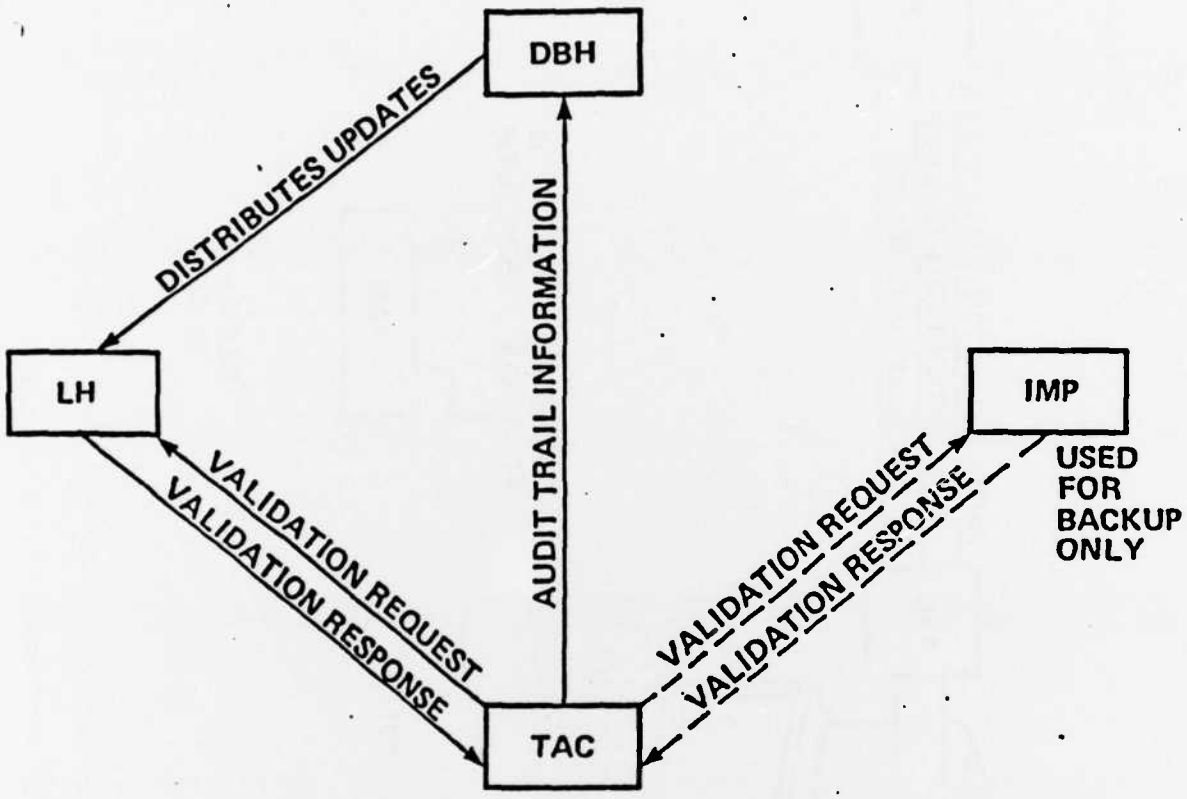

FIGURE III-21. Phase II TAC Access Control System

into the Smartcard, which will calculate a response which the user will send to the NAS. If the correct response is presented, the user will be authorized to utilize the network. This system is very strong, relying upon the user having a unique physical device and a specific password in combination with the randomly-generated challenge string.

Many of these techniques may be useful to the WIS. The idea of the self-authenticating password may be of practical application immediately to provide a network-wide user authentication process. The SAP calculation could be carried out in the LAN BIU with the information supplied either by the work-

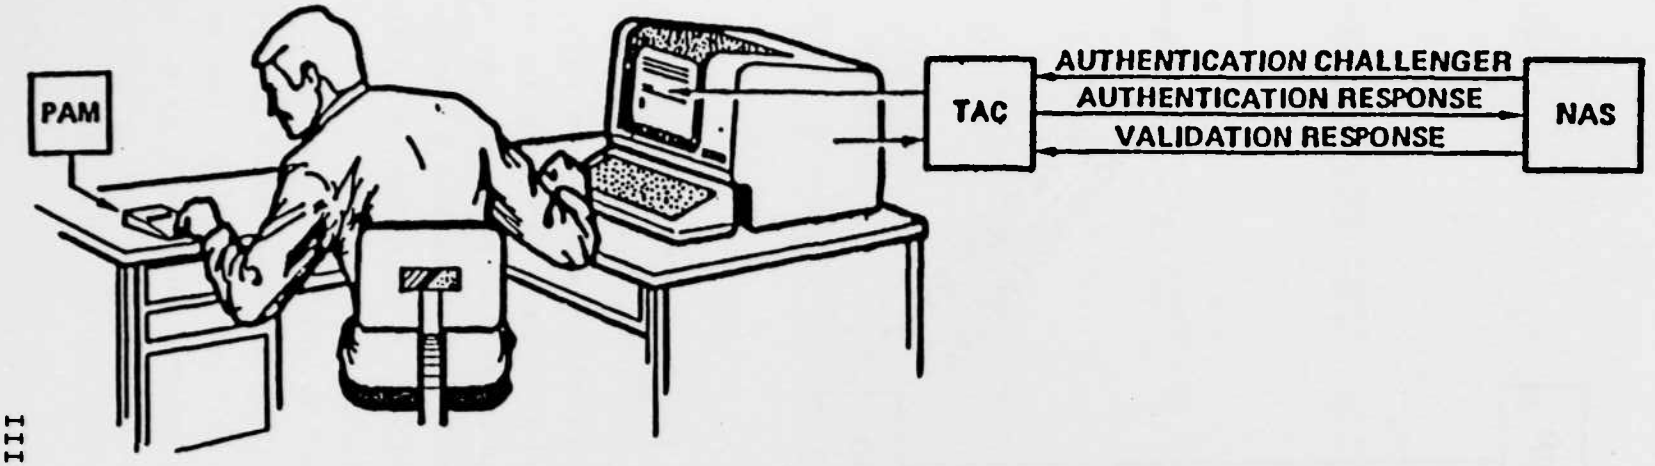

09-III

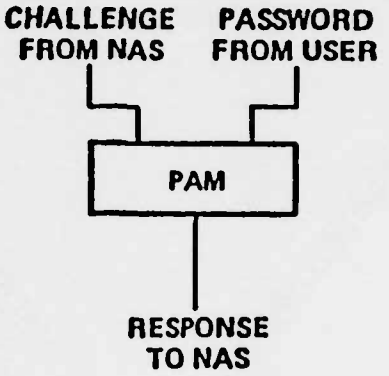

USER TYPES AUTHENTICATION RESPONSE COMPUTED BY PAM AS A CRYPTOGRAPHIC FUNCTION OF:

- **AUTHENTICATION CHALLENGE FROM NAS**
- **PASSWORD KNOWN BY USER**
- **PAS - SPECIFIC KEY**

FIGURE III-22. Personal Authentication System

to allowing the workstation to initiate a new access to the network. Either of these approaches provide an attractive near-term alternative to the prospect of a long delay in the availability of E3 or trusted LAN components. These same user authentication approaches can be retained once either E3 or trusted LANs become available.

6. Trusted Host Environment

The previous discussion focused on WIS LANs where only untrusted hosts were attached to the net. It is a goal of the WIS program that trusted hosts be introduced as soon as possible so that more effective processing of several levels of sensitive information can be performed. There are several implications on the user authentication process when trusted computers are introduced into the network. Figure III-23 shows a typical two-site WIS system with trusted hosts attached.

As described in Section IV, this structure requires that either the network and all its attached devices must operate at a network system-high level or one must have some form of trusted LAN or E3 system in place. In the first instance the Message Authentication Check must be used to ensure the integrity of all messages passed over the network. All untrusted components must operate at the network system-high. The second approach allows untrusted devices to operate below network system-high since the trusted LAN ensures the enforcement of the network security policy.

With either approach, untrusted workstations must have a means of being purged of all previous security sensitive data prior to being attached to a process at a lower security level. This is the classic periods processing purge problem and the ability to perform this action efficiently should be a significant factor in the selection of the WIS workstation.

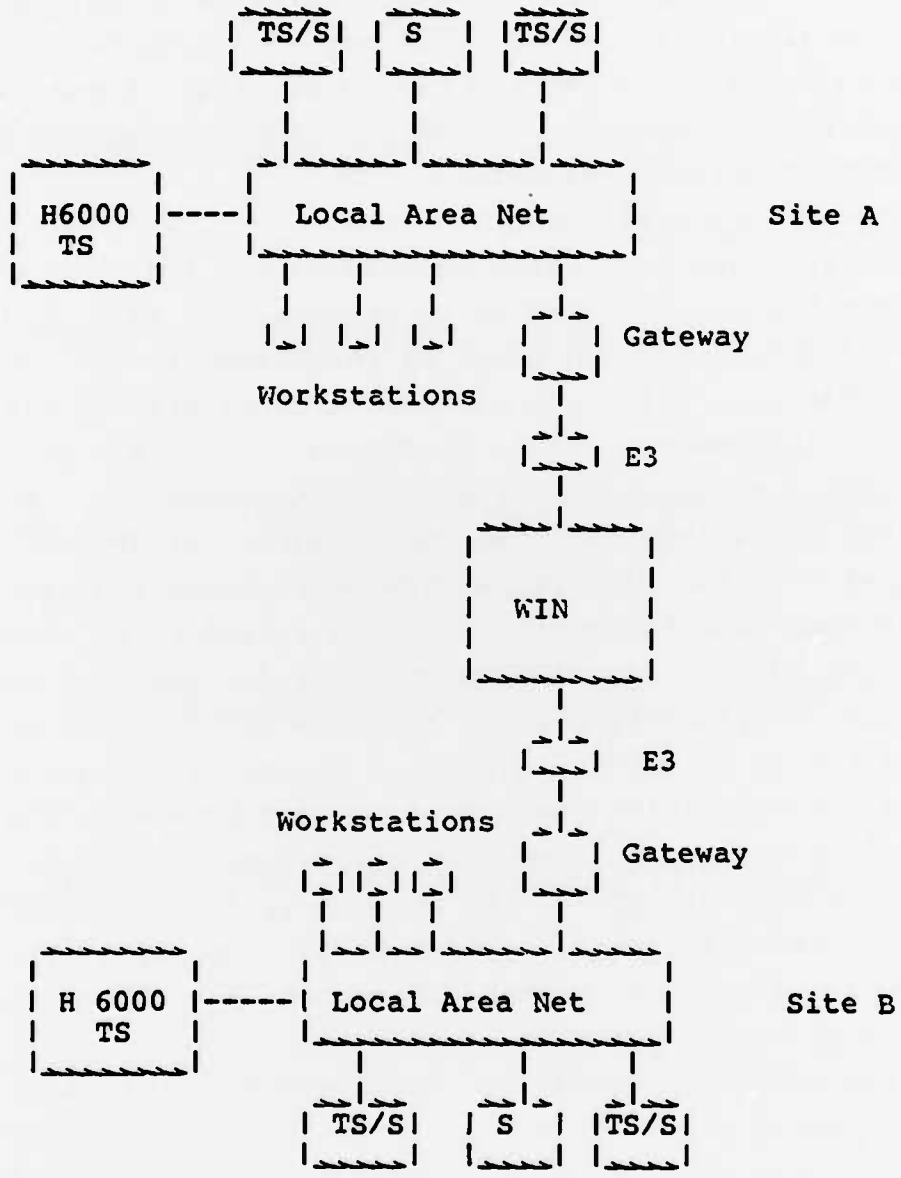

FIGURE III-23. WIS LAN Environment with Trusted LAN, Trusted and Untrusted Hosts

REFERENCES

1. William H. Blankertz, David A. Gomberg, "WIS Local Area Network Issues," Mitre Report MTR-82W00123, July, 1982.

2. Robert W. Shirey, "FY81 Final Report: Cable Bus Applications in Command Centers, Security Issues," Mitre Report MTR-81W00248-02, February 1982.

3. Deepinder P. Sidhu, Morrie Gasser, "Design for a Multilevel Secure Local Area Network," Mitre Report MTR 8702, March 1982.

4. Willis Ware, "Security Controls for Computer Systems, Report of the Defense Science Board Task Force on Computer Security," Rand Corp LR-609-1, originally published February 1970, reissued October 1979.

END

FILMED

5-85

DTIC